

CARNEGIE MELLON UNIVERSITY

School of Architecture

College of Fine Arts

Thesis

Submitted in Partial Fulfillment of the requirements for the degree of

Master of Science in Computational Design

TITLE:

Building a Bidirectional Bridge Between the Digital and
Physical Worlds

AUTHOR:

Violet Yinuo Han

ACCEPTED BY ADVISORY COMMITTEE:



Professor Eddy Man Kim

Principal Advisor

May 9, 2023

DATE

Professor Alexandra Ion

Advisor

May 8, 2023

DATE



Professor David Lindlbauer

Advisor

May 10, 2023

DATE

Acknowledgements

I'm thankful for my advisors Alexandra Ion and David Lindlbauer for leading me into research, their day-to-day support and guidance, and just being the cool (edgy) researchers they are. I'm thankful for my advisor Eddy Man Kim for helping me develop this thesis, and the insightful conversations we share.

All my friends have made my master's years both fun and fulfilling. Here I especially would like to acknowledge Hyunsung Cho, Kiyosu Maeda, Yukang Yan, Tomas Cabezon Pedroso, Daehwa Kim, Wesley Deng, and Yiming Zhang, who are with me through the fun times and down times, and constantly inspire me with who they are.

My undergraduate mentors Franc Nunoo-Quarcoo and Maria Phillips saw potential in me and helped me become a better person.

I'm thankful for my parents who never asked me to be anyone but myself, and always encouraged me to define my own path.

Abstract

We humans today are dual-citizens inhabiting both the digital and physical worlds; Both offer us unique advantages. The digital world offers us advantages such as unlimited access to information and expansion of human capabilities, while the physical world has long been providing us structural affordances, functional mechanisms, rich sensations, and much more. Currently, there exists a gap between the two. We constantly find ourselves having to frequently switch contexts to interact with either. Within this thesis, I propose building bi-direction interactions between the two worlds so that we can leverage advantages of both worlds simultaneously. I demonstrate this vision with an interactive system and a device each seamlessly bring advantages of one world into another, leveraging computational techniques and inspirations from design.

Contents

Acknowledgements	2
Abstract	3
1 Introduction	9
2 Background	11
3 BlendMR: A Computational Method To Integrate 2D Applications into Physical Objects Using Mixed Reality	12
3.1 Introduction	12
3.1.1 Contribution	15
3.2 Related Work	15
3.2.1 Blending Physical and Virtual	15
3.2.2 Adaptive MR Interfaces	16
3.2.3 Texture Mapping	18
3.3 Implementation	18
3.3.1 Input	19
3.3.2 Overview	21
3.3.3 Optimization	22
3.3.4 Implementation Details	25
3.3.5 Influence of Individual Optimization Parameters	26
3.4 Evaluation	26
3.4.1 Participants and Apparatus	27
3.4.2 Study Design	28
3.4.3 Results	31
3.4.4 Discussion	33
3.5 Application and Additional Results	34
3.5.1 Application	34
3.5.2 Additional Results	35
3.6 Discussion	36
4 Parametric Haptics: Versatile Geometry-based Tactile Feedback Devices	39
4.1 Introduction	39
4.1.1 Parametric Haptics	40

4.1.2	Contribution	42
4.2	Related Work	42
4.2.1	Tactile Perception	43
4.2.2	Portable and Wearable Haptic Feedback Devices	44
4.2.3	Embedding Functionality Within Geometry	45
4.3	Designing Parametric Haptics	45
4.3.1	Actuation Geometry	46
4.3.2	Fabrication	47
4.3.3	Design Tool	48
4.3.4	Motorized Actuation Interface	49
4.4	Design Space	50
4.5	User Study: Perception of Geometric Parameters	51
4.5.1	Study Design	52
4.5.2	Participants and Apparatus	52
4.5.3	Procedure	54
4.6	Results	55
4.6.1	Quantitative Rating of Stimuli	56
4.6.2	Qualitative feedback	57
4.7	Application Examples	59
4.7.1	Haptic feedback in Virtual Reality	60
4.7.2	Augmenting objects	60
4.7.3	Wearable notification device	61
4.8	Discussion	63
5	Discussion and Conclusion	64

List of Figures

3-1	We present an automatic approach for integrating virtual 2D applications into physical objects. Our approach takes the geometry of a physical object and an annotated 2D application as input, and produces a mapping of important sub-elements of the application to the geometry. The result can then be viewed in Mixed Reality. <i>Right</i> shows live view of the results as shown in an Varjo XR-3 headset.	13
3-2	Example of applications from Rico dataset annotated with importance scores. The bounds are retrieved from the RICO dataset. Numbers indicate the importance of individual elements (1 low, 5 high), as decided by an annotator.	19
3-3	Overview of the input processing for the geometry of physical objects. . . .	20
3-4	Overview of the system workflow, from input (geometry of physical object and 2D application) to the output of the optimization and the MR version of the output. <i>AR output is recorded live through the Varjo XR-3 headset.</i> .	21
3-5	Results of varying the different sub-objectives of the optimization.	27
3-6	Participant performing task in the two study scenarios.	28
3-7	Participant’s MR view during the study. <i>Left shows the 2D baseline and the BlendMR condition in the living room. Right show the interfaces displayed in the MR headset’s high-res focus display, enabling BlendMR to display interfaces elements that are legible.</i>	29
3-8	Mean subjective ratings from 1 (strongly disagree) to 5 (strongly agree). Error bars indicate standard error. Lower ratings for clutter and distraction are better.	32
3-9	Three implemented application scenarios. In the meeting scenario, the red folder and gray decorative object hold virtual objects. In the living room, weather applications and quick access for other apps are visible. Similar pairings between 2D applications and physical objects are shown in the social interaction scenario. <i>All images were recorded live through the Varjo XR-3 MR headset.</i>	34
3-10	Additional results with physical objects and applications. The last row is the live view from a MR headset (Varjo XR-3).	35
3-11	Additional results of different applications mapped on the same objects, as seen through MR headset (Varjo XR-3).	36

4-1	We present customizable and versatile haptic patches that are lightweight, reconfigurable, and can be integrated as wearables or into objects. (a) Users can customize their own haptic patches with a software aid we provide. (b) One linear pull can actuate many tactors, producing haptic sensation on a large area of the skin. (c) Our approach can be integrated as wearables, into existing small objects, and designed for VR experiences.	40
4-2	(a) Our tactile feedback device consists of custom tactile patches that can be connected to a generic actuation interface, that (b) is wearable. (c) Multiple tactors are actuated to stroke over the skin, while the tactor design is user-defined.	41
4-3	(a) An unactuated haptic patch with geometry recommendations for each part. (b) The same patch with its tactors actuated.	47
4-4	Users can input parameters of the haptic patch they want to create, an design tool automatically generates a working actuation geometry, along with a correctly angled guidance plane for users to design their own tactors.	48
4-5	Modular wearable interface. Users can exchange patches of different sizes for various use scenarios with the same wearable interface.	50
4-6	Our haptics patches have a wide design space for creating different tactile sensations.	50
4-7	Independent variables are reflected in tactors and their layouts. We tested 18 patches in total.	53
4-8	Study setup. Before conducting the study, each participant's arm is first calibrated for consistent indent of tactors into the skin.	54
4-9	Mean subjective ratings from the experimental questionnaires. Error bars indicate standard errors, (*) indicates statistical significant differences. . .	55
4-10	(a) A haptic patch for a VR jungle expedition game (b) a user being swept by leaves in the VR game, feeling the sensation on their forearm (c) a user being bitten by spiders in the VR game, with the haptic patch rendering this sensation.	59
4-11	(a) A row of 3 tactors integrated into an existing pair of glasses (b) a user wearing a pair of Parametric Haptic-integrated glasses in daily life, the glasses can unobtrusively and effectively deliver information.	61
4-12	(a) This back wearable patch can deliver various notifications and information to the user's lower back. (b) The user gets a gentle notification for a text message from their family. (c) We show the various tactors and directions we can render on this larger area.	62

List of Tables

3.1	Weights of the sub-objectives used parameter exploration.	26
-----	---	----

Chapter 1

Introduction

Both the digital and physical worlds offers us unique advantages. On the one hand, the digital world brings us new capabilities, access to unlimited information, and so on. On the other hand, the physical world has long been offering us structural affordances, functional mechanisms, rich sensations, and much more. We humans today are dual citizens of both worlds, interacting with and leveraging advantages of the two worlds in numerous scenarios for both work and leisure, through day and night.

However, there currently exists a gap between the two worlds. We are unable to access domain advantages brought about by the digital and physical worlds at once, and find ourselves constantly having to switch contexts throughout our day to interact with either.

I propose the creation of a perceptually hybrid space between the two worlds, with bidirectional interacterables allowing people to simultaneously access advantages of both worlds.

Within the scope of this thesis, I present two research projects each bringing domain advantage from one world to another. In Chapter 3, I illustrate a research effort to bring the unlimited access of information from the digital world into the physical world in an unobtrusive manner; In Chapter 4, I demonstrate another research that renders rich sensations of the physical world for the digital world. These two research are my initial endeavors towards bridging the digital and physical within my master's studies, building on forerun-

ners' ideas and efforts, I aim to continue such effort of bridging these two worlds with my upcoming research career.

Chapter 2

Background

Interfacing between the digital and physical has been undoubtedly an well-sought among practitioners in fields related to designing for and human factors around computing. One of the most famous paradigms on this vision was proposed by Hiroshii Ishii, composed of two research agendas he named "tangible bits" [59] and "radical atoms" [58]. Within "tangible bits", he and his group proposed providing physical forms to digital interfaces to bring computing interfaces beyond pixels [56], giving users physical control [119], taking advantage of human's long-developed capabilities of grasping and manipulating objects from the physical world [57]; Within "radical atoms", efforts have been geared towards transformable materials [58, 95], shape displays that can change its shape to display various information as well as interact with users [39], and so on. Some other works related to this vision have explored editing on a physical object and simultaneously reflecting the edit on a software interface [109]; Giving a physical object digital qualities such as being disappearable [94], as if how one could close or hide a user interface; Using semantic associations between digital contents and their physical surrounding to more-meaningfully display information in Mixed Reality [21]; Using new capabilities brought about by the digital domain to generate and fabricate new actuation mechanisms for existing physical objects [5].

Chapter 3

BlendMR: A Computational Method To Integrate 2D Applications into Physical Objects Using Mixed Reality

The following chapter is based on *BlendMR: A Computational Method To Integrate 2D Applications into Physical Objects Using Mixed Reality*, submitted to *The ACM Symposium on User Interface Software and Technology 2023*.

3.1 Introduction

Mixed Reality (MR) has the potential to change how users interact with digital information. Currently, user interface elements are presented on the displays of desktop computers, smartphones, or tablets. With MR, and especially with future always-on head-mounted displays (HMDs), personalized virtual contents can be displayed freely in space, and contain nearly arbitrary amounts of information. This enables more flexible access to information,




Figure 3-1: We present an automatic approach for integrating virtual 2D applications into physical objects. Our approach takes the geometry of a physical object and an annotated 2D application as input, and produces a mapping of important sub-elements of the application to the geometry. The result can then be viewed in Mixed Reality. *Right* shows live view of the results as shown in an Varjo XR-3 headset.

since users can integrate information directly into their environment, similar to current ambient displays and decorative objects like clocks or digital picture frames.

Current solutions to MR, however, largely focus on displaying 2-dimensional windows in specific areas of users’ environment (e.g., [21, 20, 28, 84, 80]). Their placement relies on factors such as the geometry in an environment [41], or the semantic connection between the physical and virtual objects [21]. While this simplifies access to information, adding multiple applications in an environment can lead to clutter and distraction. Current MR interfaces lack the unobtrusive qualities of peripheral objects and ambient displays. This challenge is exacerbated if MR systems aim at enabling continuous access to the dozens of applications an average smartphone user has installed [43]. While providing users with menus to toggle the visibility of individual applications and contents could alleviate this challenge, this approach re-introduces the need for users to explicitly search for the information they need. This eliminates the key benefit of simple and unobtrusive access to information offered by MR.

To address this challenge, we present BlendMR, a novel approach that embeds information from existing 2D applications into physical 3D objects. BlendMR takes as input a 2D smartphone application, specifically app screenshots and information about interface hierarchy and element importance from the Rico dataset [26]. Additionally, BlendMR takes 3D scans of the objects in a room as input. Users then decide what physical object an application should be mapped onto, e.g., the weather app onto a physical air purifier as shown in Figure 4-1.

Our novel approach analyzes the physical object, and finds areas that are suitable to “host” virtual elements based on aspects such as curvature and position in a room with respect to users. Using combinatorial optimization, BlendMR then maps the elements of the application to the geometry of a physical object, and displays the information through a see-through head-mounted MR display. Our approach effectively transforms 2D applications into ambient information displays embedded in everyday objects. BlendMR solves the problem of how to integrate secondary virtual applications with physical objects for simple and unobtrusive information access, minimizing potential clutter and distraction. *All MR examples in this paper are recorded live through a Varjo XR-3 video see-through MR headset.*

We evaluate our approach in a user study ($N = 16$), comparing BlendMR with a standard 2D window baseline. Participants performed reading tasks in two different scenes (living room, desk) and simultaneously answered questions regarding the contents of virtual interfaces in the scenes. Results show that BlendMR statistically significantly reduced clutter and distraction, and was preferred over the baseline. We did not observe any reduction in answering speed between the two display modes. We believe that this highlights the applicability and usefulness of BlendMR. We demonstrate the potential of our approach in a series of application scenarios, specifically, casual interaction in a living room, social interaction, and an office setup. Finally, we provide insights and qualitative results on the impact of individual sub-objectives of the proposed optimization method, and additional results of 15 individual mappings.

BlendMR is a new type of MR content presentation that allows users to continuously access personalized information while reducing clutter. It showcases a potential future of MR in which virtual applications are no longer directly transferred from traditional computing devices such as smartphones, but are more tightly integrated into users’ physical environments.

3.1.1 Contribution

In summary, we make the following contributions:

- A computational approach to map information from 2D applications onto physical 3D objects, called BlendMR. This approach takes standard annotated 2D applications and the geometry of 3D physical objects as input, and uses geometry processing and combinatorial optimization to find suitable mappings.
- A comparative user study ($N = 16$) showing that BlendMR decreases clutter and distraction, and is preferred when compared to a baseline of traditional 2D MR interfaces.
- A set of results and application scenarios highlighting the versatility of the approach for future MR interfaces.

3.2 Related Work

We present related works on blending physical and virtual objects, adaptive MR interfaces, and relevant works on texture mapping, as our approach effectively maps 2D objects (e. g., images) onto 3D geometry.

3.2.1 Blending Physical and Virtual

MR naturally holds the promise of blending the physical and the virtual, with augmented contents overlaid on top of the physical world. For efficient interactions in everyday MR, virtual and physical contents need to be aligned. Prior research leveraged physical planar surfaces to host virtual contents, such as the work by Nuernberger et al. [96] or Ens et al. [34]. Geometric alignment increases the integration level, and is useful for view management and general content adaptation. With HeatSpace, Fender et al. [37] use similar approaches to adapt the placement of physical displays in a space to best suit users' behavior. In

OptiSpace [38], this approach is extended to adapt content displayed through projection mapping. They only consider users’ viewing angle, whereas a core component of BlendMR is to account for the geometry of the object that hosts the virtual elements, and to decide what parts of 2D applications to display based on their importance. Gal et al. [41] took the geometry of a space into account and used heuristic optimization to adapt virtual contents to conform to a space, e.g., by adapting a virtual race track to the furniture in a room. They did not leverage physical contents as containers for 2D applications. Approaches such as SceneCtrl [127], Oasis [110], Remixed Reality [81], Substitutional Reality (Simeone et al. [108] and Suzuki et al. [113]) and TransforMR [67] overlap virtual contents onto physical objects with varying degree of immersion. Other research in spatial augmented reality [91, 105] and tangible user interfaces [15, 51, 29] show that a spatial mapping of digital information and interactions onto physical geometries provides advantages such as haptic feedback, ubiquitous access to information, and simplified interaction. These works rely on manually mapping virtual elements to physical objects. Given that content creators cannot anticipate which physical objects are in users’ space, or what virtual elements they would like to use, these approaches do not scale outside of a lab environment. We complement these approaches by providing a computational approach to map virtual applications and physical objects in everyday MR environments with the goal of simple access and decreased distraction. BlendMR moves beyond 2D windows towards a tighter integration of virtual contents with physical objects, shown by our evaluation to be beneficial in terms of clutter, distraction, and user preference.

3.2.2 Adaptive MR Interfaces

Existing research explored approaches to adapt virtual interfaces for labeling or annotation placement problems [9, 33, 103]. Azuma and Furmanski [6] evaluated their cluster-based label placement algorithm and found that it was the best in terms of objective and subjective measures among four different algorithms. Gebhardt et al. [42] proposed a system that adapts virtual labels based on eye movement data by employing reinforcement learning. We leverage these insights for creating the connection between virtual applications and

physical objects, and for controlling the visibility of applications. Our approach then creates a representation that optimally integrates elements of the virtual applications into physical objects.

Other research has taken environmental factors such as physical objects and geometry into account, such as the work by et al. [72], Nuernberger et al. [96], Fender et al. [38], or Tahara et al. [114]. Ens et al. [35] consider spatial consistency, surface structure and visual saliency of surrounding environments to manage interface layouts in VR. Matulic et al. [92] proposed a method that automatically detects suitable physical surface areas for projecting digital information. SemanticAdapt [21] adapts virtual interface layouts by considering the semantic relationships between physical and virtual objects. They leverage insights from computational interaction (cf. Oulasvirta et al. [98, 99]) and adaptive user interfaces (e.g., Supple [40], ADAM[101]) for to generate and adapt MR layouts.

Lu et al. suggested using the users’ periphery to place and access secondary information [85, 84]. Virtual elements are head-anchored 2D windows, which leads to visual clutter if too many elements are displayed. We provide a way to reduce clutter by embedding virtual elements into physical objects.

Lindlbauer et al. [80] used combinatorial optimization (cf. [97]) to control the placement and appearance of virtual applications based on users’ cognitive load, their task and environment. We similarly solve an assignment problem, and provide a novel formulation that takes the geometry of physical objects into account. The importance of surrounding environments for placing virtual elements is also supported by work such as by Luo et al. [87], who considered aspects of the physical environment for placement of objects. DiVerdi et al. [28] generated different levels of details for interfaces to fit users’ needs. Our work takes a similar approach by taking importance into account for integrating virtual elements into physical objects.

We extend prior work by not only taking the environment and physical objects into account for placement, but by considering their surface geometry for integrating virtual elements. To achieve this, our method leverages geometry analysis and an optimization-based method to

detect suitable sub-surface areas and adapt virtual content layouts, respectively, to fit those surfaces.

3.2.3 Texture Mapping

Our approach is inspired by general texture mapping methods [90, 50]. Texture mapping refers to the creation of a faithful mapping of 2D image data onto 3D objects. This typically involves parameterization and matching. There exist various approaches in the computer graphics community involving methods such as image warping [107], mesh warping [89], or methods that use using neural textures [125]. We leverage standard methods for parameterization, mapping and texture packing [77]. In contrast to work in traditional texture mapping, we do not wish to achieve full coverage, but to sparsely add elements of 2D applications onto the geometry of existing physical objects. We model this as linear programming problem, which can be solved efficiently and in real time. We hope to expand our work and explore other gradient-based methods or learning-based methods for mapping 2D applications onto 3D objects in the future.

3.3 Implementation

The goal of our approach is to reduce clutter when displaying MR applications. To achieve this, we embed 2D applications into existing physical objects. In other words, we take the individual interface elements of a 2D application and map them as textures onto the surface of physical objects in MR. Our approach is designed to serve as a peripheral display showing important information, but that is not directly interacted with. For interaction, we envision future approaches to adapt the representation to ones that are well suited for interaction (cf. Cheng et al. [20]). Additionally, our approach avoids overlaps, follows established hierarchies, and is aligned in groups (cf. Dayama et al. [25]).

3.3.1 Input

We leverage two main sources of inputs: annotated 2D applications and 3D meshes of existing physical objects. The mapping between applications and physical objects is created by users by moving an application close to a physical object. This approach could easily be replaced with automatic placement approaches [21, 6, 9] that leverage semantics, physical proximity, or host characteristics. BlendMR then decides *how* applications should be blended into 3D objects based on the design of the 2D applications, the geometry of the physical objects, and the users’ common position in a room (cf. [37]).

2D applications

We use applications from the Rico dataset [26], with examples shown in Figure 3-2. The dataset consists of mobile phone applications. Each application is divided into individual interface elements, denoted as E in our context, and their bounding boxes. We additionally extract information if pairs of elements are adjacent, denoted as $r_{e_i, e_j} \in \{0, 1\}$. We consider the distance between elements with respect to the overall size of the application, and whether they are horizontally or vertically aligned, which usually encodes grouping relationships.




Figure 3-2: Example of applications from Rico dataset annotated with importance scores. The bounds are retrieved from the RICO dataset. Numbers indicate the importance of individual elements (1 low, 5 high), as decided by an annotator.

We extend each element with information on its importance within the context of an application. Figure 3-2 illustrates an example where an annotator added importance values




Figure 3-3: Overview of the input processing for the geometry of physical objects.

where the individual elements would be used as ambient display in a room. This information is added by an annotator as values between 1 (not important) and 5 (highly important) for each interface element of an application using a custom annotation tool. The information is normalized to the maximum importance value (5), and stored as importance values $i_e \in \{0, 1\}$. We chose a manual labeling process, as we focus on finding an optimal mapping between a physical object and a 2D application.

Extending this with task-specific consideration, or to further automate the process (e.g., Wu et al. [124], Chen et al. [19]) provides an interesting direction for future work. The size of the elements is decided before the optimization, with a pre-determined scale factor for all elements based on the size of the target object, and a minimum size constraint to ensure the readability of the displayed output.

Geometry of physical objects

To integrate virtual applications into physical objects, our system requires mesh representations of physical objects in users' environment (see Figure 3-3). We leverage 3D scans of physical objects, and store the representations as standard meshes with vertices and faces, denoted as $M = \{V, F\}$. We first segment the mesh into individual segments S to simplify parameterization. Each segment can host multiple elements of an application. We use spectral clustering [82] for segmentation, which uses a combination of geodesic and angular distances between segments. We then parameterize each segment to retrieve the texture maps T using least squares conformal maps [77]. The maps are scaled according to their relative size on the overall mesh M and aligned to the global horizontal and vertical axis by orienting them so that a randomly chosen 2D horizontal line on the parameterization aligns




Figure 3-4: Overview of the system workflow, from input (geometry of physical object and 2D application) to the output of the optimization and the MR version of the output. *AR output is recorded live through the Varjo XR-3 headset.*

with a 3D horizontal line on the mesh. Finally, we create a grid of placement candidate slots for each texture, denoted as G , by overlaying a grid on top of the texture maps. We place all individual textures within a single rectangle using a rectangle-packing solver (cf. Sander et al. [106]), and overlay it with an equally spaced grid (see Figure 3-3 *center*). Grid cells are oriented correctly as the texture maps are axis aligned. Each grid cell contains multiple vertices and faces, depending on the mesh resolution.

3.3.2 Overview

In the following, we describe how we map the individual interface elements of one application onto a physical object. A system overview is illustrated in Figure 3-4. Our system analyzes the geometry of the physical objects, specifically its segments, curvature, distortion in the parameterization, and position with respect to users. This information is used to generate a set of placement candidate positions on the mesh. We then utilize integer programming to decide which elements of an application to place, and where on the mesh. To achieve this, our approach considers objects' geometry, the importance of virtual elements, and their relation and hierarchy in the 2D application.

3.3.3 Optimization

We map the elements of a 2D application e onto the geometry of a physical object, specifically the grid cells g on the texture map as retrieved from the mesh segmentation and parameterization. Each element can cover multiple grid cells. This decision is formulated as common in assignment problems as

$$x_{e,g} = \begin{cases} 1 & \text{(if } e \text{ is assigned to } g) \\ 0 & \text{(otherwise)} \end{cases} \quad (3.1)$$

We formulate the overall objective as a weighted sum of multiple individual sub-objectives. The sub-objectives are informed by general usability guidelines, and prior work on adaptive user interfaces [25, 21, 80, 101]. Specifically, we prioritize the presentation of important objects (F_{imp}), ensure that the hierarchies of the 2D applications remain intact, and that related elements are grouped together (F_{dis}). Additionally, we ensure that the areas on the physical objects are well suited for "hosting" elements, i. e., they are well visible, have low curvature, and do not contain features such as wrinkles or tightly packed triangles which could lead to distortion (F_{geo}). Lastly, we optimize for users' primary viewing direction (F_{view}). We formulate the main objective with the considerations as

$$\underset{e,g}{\operatorname{argmin}} w_{\text{geo}} \cdot F_{\text{geo}} + w_{\text{view}} \cdot F_{\text{view}} + w_{\text{imp}} \cdot F_{\text{imp}} + w_{\text{dis}} \cdot F_{\text{dis}}. \quad (3.2)$$

In the following, we detail the individual sub-objectives. In addition to those, we leverage the overlap-avoidance from Dayman et al. [25].

Geometry

Each grid slot g hosts multiple vertices and faces. For each grid slot, we compute the average Gaussian curvature $K(g)$ of the containing vertices. Additionally, we count the number of vertices that fall within a cell $\gamma(g)$, normalized for the minimum and maximum number of vertices in all cells per segment. Our approach penalizes high curvature and high number

of vertices in a cell, as this leads to visual distortion, formulated as

$$F_{\text{geo}} = \sum_e \sum_g x_{e,g} \cdot \frac{1}{2} (K(g) + \gamma(g)). \quad (3.3)$$

Viewing direction

BlendMR penalize if grid cells are not facing users, as this decreases the visibility of visual elements. We calculate the dot product between all normals in a cell ($\mathbf{n}_g$) the vector from users to the cell ($\mathbf{d}_{u,g}$), and average the results, denoted as $\alpha(\mathbf{n}_g, \mathbf{d}_{u,g})$. The sub-objective is formulated as

$$F_{\text{view}} = \sum_e \sum_g x_{e,g} \cdot \alpha(\mathbf{n}_g, \mathbf{d}_{u,g}). \quad (3.4)$$

Importance

This sub-objective penalizes if important elements are not displayed, formulated as

$$F_{\text{imp}} = \sum_e \sum_g (1.0 - x_{e,g} i_e). \quad (3.5)$$

Distance

We aim to keep elements that are adjacent in the 2D application close together when placed on the geometry of the physical objects. This is reflected in the distance sub-objective as

$$F_{\text{dis}} = \sum_{e_i} \sum_{e_j} \sum_g x_{e,g} \cdot r_{e_i, e_j} \cdot d_{e_i, e_j}. \quad (3.6)$$

r_{e_i, e_j} specifies if two elements are directly adjacent in the 2D application. We then minimize the distance for these elements. We chose this approach to minimize computational complexity by avoiding a full-factorial quadratic formulation. For elements that are directly adjacent, we minimize their distance d_{e_i, e_j} in the grid, calculated as follows. Let c_i, c_j be centers of elements e_i and e_j . We use $\theta = \arctan \frac{(c_i - c_j)_x}{(c_i - c_j)_y}$ to denote the angle between the

vector connecting the centers of the pair and the x -axis. In addition to encouraging pairs in close proximity, we further penalize pairs that are not along any of the four cardinal directions (i. e., N, S, E, W). We calculate the distance as

$$d_{e_i, e_j} = \|c_i - c_j\|_2 \cdot (2 - \cos(4\theta)).$$

We chose this formulation as this corresponds to typical 2D layouts that exhibit a certain degree of regularity.

Constraints

While the above objective produces potentially meaningful results, we introduce additional constraints to avoid trivial solutions and preserve information from the initial 2D application.

As a first constraint, similar to prior work on adaptive interfaces [80, 21, 101], we ensure to avoid duplicating virtual elements in the layout. Additionally, we enforce a certain level of layout consistency. If an element is above another one, for example, it should not be placed below it when placed onto the physical object, but remain on top, or to its left or right. The constraints enforce similar consistency for the other three directions. To achieve this, we integrate a prior formulation introduced by Dayama et al. [25]. We introduce the same binary variables indicating the relative position between two elements e_i and e_j , specifically if e_j is above e_i (Γ_{e_i, e_j}) or left of e_i (Π_{e_i, e_j}). We refer readers to Dayama et al. [25] for details on how those two variables are constructed. We then construct four constraints, one to constrain the vertical placement

$$\Gamma_{e_i, e_j} = 0 \quad (\text{if } e_i \text{ is above of } e_j) \quad \forall e_i, e_j \quad (3.7)$$

and one to constrain the horizontal placement

$$\Pi_{e_i, e_j} = 0 \quad (\text{if } e_i \text{ is left of } e_j) \quad \forall e_i, e_j. \quad (3.8)$$

Note that in this formulation, for any element pair e_i, e_j , we prohibit that they are above each other if e_j is below e_i . Finally, we use the approach of Dayama et al. [25] to avoid overlapping elements.

3.3.4 Implementation Details

The optimization is implemented in Python 3.8 using Gurobi 9.5 [45] for solving the constraint optimization problem. We use libIGL[61] for mesh processing, including generation of textures and calculation of Gaussian curvature. The processing takes between 30 and 120 seconds on a Macbook Air M1 2020. The geometry of the physical objects used in this paper was scanned using an EinScan Pro 2X Plus handheld industrial scanner. For displaying objects in AR, we use the Varjo XR-3 headset connected to a commodity gaming computer (Alienware R12). The device is a high-resolution video see-through display that enables overlaying virtual contents on physical objects without having to consider radiometric interference. The 2D applications are displayed as textures onto manually aligned virtual representations of the physical objects. In the future, we hope to extend our approach with automatic digitization and tracking of 3D objects in an environment. We chose static 3D scans to evaluate the feasibility of our approach.

Computational considerations

Our approach contains quadratic terms for the relation between elements, negatively impacting the run-time. To address this challenge, we only take distance terms into account for elements that are directly adjacent. Additionally, we limit the number of possible grid slots to typically between 100 and 200, depending on the size of the mesh. Lastly, we decide on the size of the elements based on their importance before the optimization, to avoid this quadratic term. In practice, this decreases the computational complexity considerably, allowing for a fast convergence time on a commodity laptop. We believe this is sufficient, as we do not anticipate our approach to run at interactive rates. Mappings between 2D

applications and physical interfaces should update infrequently to enable users to leverage their spatial memory and familiarize themselves with the integrated representation.

3.3.5 Influence of Individual Optimization Parameters

We explored the influence of the main sub-objectives for geometry F_{geo} and distance F_{dis} on performance and quality of the integration. The different weights are shown in Table 3.1. We chose to always include the sub-objective for importance F_{imp} , as its removal would lead to a trivial (empty) solution. Additionally, we do not vary the sub-objective for viewing angle F_{view} , as this would lead to random placement of elements. We compute the results with these weights for three different 2D applications, each mapped onto the geometry of a different physical object, and add results with all sub-objectives enabled. All experiments are performed on a MacBook Air M1 2020 (8 cores, 8GB Ram), and take between 0.5 and 20 seconds to compute, depending on the mesh resolution and optimization parameters, with higher speed for fewer parameters. Figure 3-5 shows the results of this exploration. Without

ID	w_{dis}	w_{geo}	w_{imp}	w_{view}
All sub-objectives	0.25	0.25	0.25	0.25
No distance	0.0	0.33	0.33	0.33
No geometry	0.33	0.0	0.33	0.33
Importance only	0	0.0	0.5	0.5

Table 3.1: Weights of the sub-objectives used parameter exploration.

the distance term, elements are generally further apart, as expected, breaking the hierarchy of the 2D application. Without the geometry term, elements are placed in areas with high distortion (e.g., front part of bunny; bottom edge for box). Using only the sub-objective for importance exacerbates those effects particularly apparent for the speaker. We believe that these results show that all constraints are valuable components of the formulation.

3.4 Evaluation

We performed a comparative evaluation ($N = 16$) to gather insights into the efficacy of BlendMR. Specifically, we were interested in users’ ability to find and read contents on




Figure 3-5: Results of varying the different sub-objectives of the optimization.

BlendMR interfaces, and BlendMR’s ability to reduce perceived clutter and distraction. Participants were asked to read a newspaper article on paper. At the same time, they were asked questions by an experimenter, which could be answered using virtual interfaces displayed through a video see-through MR headset. As representation of virtual contents, we compared BlendMR with a baseline approach where applications were displayed as floating 2D windows, which reflects current standard MR interfaces. All blended interfaces were direct outputs of our approach, without manual adjustment. We measure the time to complete the question-answering task, error rate, and subjective ratings for preference, ease of use, clutter, and distraction. Results show that BlendMR significantly reduces visual clutter and distractions induced by virtual user interfaces, without increasing search time or error rates compared to the baseline. Participants preferred BlendMR, largely owing to the integration of the virtual and physical and the opportunity for functionality mapping.

3.4.1 Participants and Apparatus

We recruited 16 paid participants (age $M = 26.9$ years, $SD = 2.7$; 9 female, 8 male) from a local university. Participants had an average experience of $M = 2.7$ ($SD = 0.77$) in using AR interfaces and $M = 2.8$ ($SD = 0.8$) for using VR interfaces, on a scale from 1 (none) to 5

(expert). All participants had corrected or corrected-to-normal vision based on self-reports. The study was approved by the local institutional review board. The study took place in a quiet experimental room, shown in Figure 3-6. The room was divided to simulate a living room and a work station scenario. The software was implemented using Unity 2019. The baseline interfaces were applications from the Rico dataset. The blended interfaces were the same application processed using BlendMR, but with slightly modified contents (e.g., different temperature on the weather app) across conditions to avoid answering based on memory. Participants viewed the virtual content using the Varjo XR-3 video see-through headset, as shown in Figure 3-7.



Figure 3-6: Participant performing task in the two study scenarios.

3.4.2 Study Design

We use a within-subject design with two independent variables, *display mode* with two levels (baseline, BlendMR) and *environment* with two levels (living room, work station), yielding four conditions total. In each condition, participants were asked to perform a primary reading task, while occasionally answering questions regarding peripheral virtual displays, described below.

Tasks

We use a dual-task paradigm as employed by previous research in adaptive MR interfaces [80, 24, 21] to simulate scenarios where users are engaged in their environments, and the virtual user interfaces function as peripheral information displays.

Primary task As primary reading tasks, participants were given short texts printed on paper, and had to answer comprehension questions afterward. We did not measure any performance metrics for primary tasks, as they were used only as stand-ins to create more realistic scenarios of virtual user interfaces as peripheral displays in everyday living environments.

Secondary task and applications. As secondary tasks, the experimenter asked participants a series of questions that could be answered using the virtual interfaces that were shown using the baseline or BlendMR. For each condition, participants were asked seven questions, covering all applications. Questions were asked during random intervals between 30 and 60 seconds. Example questions include: "What is the current temperature?" and "What is a recent sporting event you can attend?". We selected two sets of applications




Figure 3-7: Participant's MR view during the study. *Left shows the 2D baseline and the BlendMR condition in the living room. Right show the interfaces displayed in the MR headset's high-res focus display, enabling BlendMR to display interfaces elements that are legible.*

for the two environments, including seven applications for the living room (weather, sports, microphone, movies, health, newspaper, recipe), and five applications for the work station (printshare, pizza deliver, slideshare, plane tickets, market place). The applications and mapping for the living room, as well as close-up views from the high-resolution focus

display of the MR headset are shown in Figure 3-7. Mappings between applications and objects were manually chosen based on semantic connection between the application and the objects, or because the objects were well-suited as virtual containers. Participants saw the same set of interfaces across the two display modes within each environment. To avoid that participants memorize contents, we changed the contents of the applications, while maintaining their size and appearance. For example, the weather application displayed 50 degrees in the BlendMR condition, but displayed 47 degrees in the baseline condition. Note that we are mostly interested in users' ability to read the text and perceived clutter and distraction, rather than search efficiency. BlendMR is designed for scenarios where users already know where to find the application in their environment, similar to a physical clock or decorative object.

Procedure

Participants were briefly introduced to the experiment setup. They then completed demographic and motion-sickness questionnaires, and signed a consent form. Upon the start of each condition, the experimenter briefly walked participants through the applications. This helped them to familiarize with applications placements, but avoided content memorization. Participants completed the tasks for both display modes in one environment, and then for the other environment. The order of the environment and display mode within an environment were counterbalanced using a Latin square. We did not fully counterbalance the presentation sequence of conditions across two environments, to avoid requiring participants to switch between environment setups multiples times. Each condition ended with a post-condition questionnaire. A post-study interview is taken place upon completion of all conditions. The study took approximately 45 minutes per participant.

Measurements

Secondary task competition time, i.e., time to verbally answer a question, was recorded by the experimenter using a stopwatch. Time was measured from when the experimenter

completed asking the question to when participants answered. Errors in the answers were manually recorded by the experimenter. We employed a post-condition questionnaire based on the System Usability Scale (SUS) [16] (questions on confidence, ease of use, and whether participants would like to use the tested system). We added targeted questions on information access ("Information on virtual interfaces were accessible when I needed to retrieve them."), visual clutter ("The virtual user interfaces induced visual clutter in my environment.") and visual distraction ("The virtual user interfaces induced visual distraction to my reading task."). Participants also participated in a semi-structured post-study interview, expanding on their reasoning and subjective experiences.

3.4.3 Results

Results indicate that BlendMR significantly decreases visual clutter and distraction, while maintaining perceived interactivity and legibility of the interfaces. Results did not show any statistically significant difference for the time it took participants to answer the questions between the baseline and BlendMR. Participants also provided further insights on the BlendMR system and how they would like to incorporate it in their lives with future everyday MR. Statistical analysis was performed using JASP 0.16.3 [62].

Time and error

All participants answered all questions across 4 conditions correctly. We performed a 2 (display mode) $\times$ 2 (environment) repeated measures ANOVA to gather insights into secondary task completion time. Results did not yield a statistically significant difference for display mode ($p = .425$) between BlendMR ($M = 6.09$ sec, $SD = 4.26$) and baseline ($M = 6.39$ sec, $SD = 3.89$). While there are inevitable imprecisions in the recorded times due to the manual process, we do not believe that this changed the results in a way that would favor either condition. Results yielded a main effect for environment ($F_{1,444} = 43.10$, $p < .001$), with participants answering questions faster for work scenario ($M = 5.03$ sec, $SD = 3.23$) than the living room ($M = 7.44$ sec, $SD = 4.47$). We believe that this is due to the different

object arrangements between rooms, highlighting the variability between the two setups. The results indicate that differences between BlendMR and the baseline in secondary task competition times are likely due to noise rather than systematic challenges in the legibility or overall presentation. This is reflected in participants' comments such as *"Important information are filtered through... and it doesn't cause problems to me comprehending the interfaces"* (P1); *"I think both of them (2D display system and BlendMR) are quite good (in legibility), everything was okay to read."* (P12)

Questionnaire data

We analyzed the questionnaire data using a series of Wilcoxon signed-rank tests (Bonferroni adjusted $\alpha = 0.0083$). Results of the differences between display modes are illustrated in Figure 3-8. Results indicate that BlendMR reduced perceived clutter ($Z = 3.861$, $p < .001$) and perceived distraction ($Z = 3.977$, $p < .001$) compared to the baseline in a statistically significant manner. Baseline ratings for information access were higher than for BlendMR ($Z = 2.727$, $p = .005$). Results did not yield a statistically significant difference between BlendMR and the baseline for confidence ($p = 0.139$) or ease of use ($p = 0.623$). Participants rated BlendMR higher than the baseline for the question whether they would like to use the system frequently ($Z = -2.643$, $p = .007$) Results did not show any differences across all questions for the different environments.




Figure 3-8: Mean subjective ratings from 1 (strongly disagree) to 5 (strongly agree). Error bars indicate standard error. Lower ratings for clutter and distraction are better.

3.4.4 Discussion

BlendMR successfully reduced visual clutter and distraction caused by the presence of virtual interfaces, and was ranked higher rating. This is validated by participants' comments during interviews: *"Blended interfaces induce less visual cluttering, and also less distractions. 2D interfaces are just like smartphone screens, so they are like smartphone screens floating in the air... so I think it is pretty annoying"* (P14). Both display modes were new to participants, and we aimed not to bias them towards BlendMR in our instructions.

Integration Participants preferred integrating digital information with the 3D world (*"Because I am 3D, why put everything in 2D in the 3D world"* (P12)). 2D displays induce a sense of dis-belonging and even intrusiveness to their environments, e. g., as stated by P7: *"Mapping (information onto objects) is better. This (2D display system) is very odd, they (2D interfaces) don't look like they belong"*. While all applications could simply be shown as 2D interfaces on physical planes such as tables or walls, our findings indicate an inherent value of an integrated appearance.

Information access Results show reduced questionnaire ratings for information access for BlendMR. This might be because participants are more familiar with 2D interfaces, as commented by six participants, e. g., *"2D interfaces ... are similar to smartphone interfaces ... so they are more intuitive to read"* (P14). Nevertheless, eleven participants positively commented that BlendMR's enabled efficient information access by filtering unimportant information: *"It (BlendMR) displays information in a way that highlights the most important things. In a 2D interface, (there are) some things I don't need to see"* (P13) We believe that this highlights that even though participants are used to conventional 2D layouts, they quickly saw the benefits of our novel representation.

3.5 Application and Additional Results

3.5.1 Application

We present three application scenarios highlighting the efficacy of our approach, shown in Figure 3-9 and Figure 4-1 (*right*). All applications were created using BlendMR, and displayed live in MR using the Varjo XR-3 headset. Physical objects were manually aligned, the mapping between 2D applications and the physical targets was created by a content creator.



Figure 3-9: Three implemented application scenarios. In the meeting scenario, the red folder and gray decorative object hold virtual objects. In the living room, weather applications and quick access for other apps are visible. Similar pairings between 2D applications and physical objects are shown in the social interaction scenario. *All images were recorded live through the Varjo XR-3 MR headset.*

Meeting While collaboratively working on a problem on a whiteboard (Figure 3-9, *left*), users can check status message and upcoming calendar entries, for example. While the current mapping of 2D applications to physical objects is static, we hope to expand our approach with automatic visibility control of applications users’ cognitive load and task [80].

Casual interaction in a living room A user is in their living room (Figure 3-9, *center*). Their environment contains peripheral elements for applications such as weather, calendar, mapped onto multiple decorative objects such as a metal box and a decorative house. The applications enable them to quickly check information such as the weather, unread messages, enable them to control a music application, see what movies were released.

Social interaction A user is sitting in their couch at a table, with another user on the other side of the table (Figure 3-9, *right*). On the table, there are quick-access icons for




Figure 3-10: Additional results with physical objects and applications. The last row is the live view from a MR headset (Varjo XR-3).

movies on a physical hat, as well as a peripheral weather application. The interfaces are designed to blend into the environment and fade into the background. Once users require access, the interfaces are meant to change their representation to one that simplifies input and interaction.

3.5.2 Additional Results

We further provide additional BlendMR outputs with physical objects (Figure 4-9), and different interfaces on the same objects (Figure 3-11). Our approach is able to identify suitable areas even for challenging geometry such as the desk lamp. For the red document box (2nd column), three icons that were annotated to be the most important (web pages, pictures, documents) are shown.



Figure 3-11: Additional results of different applications mapped on the same objects, as seen through MR headset (Varjo XR-3).

3.6 Discussion

We contribute BlendMR, a novel optimization-based method that maps 2D applications onto the geometry of physical objects and takes geometric constraints, and parameters of the 2D interfaces into account. The resulting mappings are used to create unobtrusive MR interfaces for peripheral usage. Current MR user interfaces typically present MR interfaces as 2D windows, either floating in space (e.g., [21, 80]) or aligned to the world (e.g., [96]). We believe that for a future where MR content is ubiquitous and continuously available, we need to reconsider how to present virtual applications to users. Few works consider scenarios where physical objects become “hosts” for elements of 2D applications. This might be because current MR applications focus on displaying task-relevant interfaces, rather than ambient information. We believe that BlendMR opens a path to explore this novel MR representation. We chose an optimization-based approach that blends 2D applications and physical objects, as we do not believe that it is feasible for content creators to anticipate what objects are in users’ surroundings. Our current implementation assumes that physical objects are generally suitable for hosting 2D applications. We do not yet take aspects such

as personal preference into account. While a decorative object might be a suitable host for 2D applications when users are working, for example, they might not want to overlay the physical object with virtual information when they are in a social context. Incorporating this type of personalization and context awareness will be interesting for future work. In general, however, based on the results of our approach and validated by the comparative evaluation, we are confident that BlendMR provides a suitable and beneficial alternative to the traditional 2D window design of MR applications.

Evaluation Our current evaluation aimed at evaluating whether BlendMR can correctly provide a representation that decreases clutter and distraction, while providing efficient access to information. In the future, we hope to explore cases where the importance scores are imperfect, and result in adjustments. This, however, we believe is partly independent of the visual integration. Similarly, we hope to expand our evaluation with more dynamic settings in the future, e.g., varying the number of physical objects and applications. This will require integrating multiple systems, e.g., for automatic placement and user tracking. Our current work has successfully show that the concept and implementation are valuable for users. Integration into a more complex system is a challenging and interesting next step that will lead to a new set of challenges for BlendMR and adaptive MR systems as a whole.

Mapping virtual contents to physical objects Our current approach only considers the shape of the physical objects, not their texture. This can lead to virtual contents overlaying important information on the physical objects. This was not an issue for the everyday objects in our tests, since many of them had uniform textures. Future versions of our approach, however, should consider the physical texture, specifically how much information it holds, as an additional parameter. Our formulation would enable a straightforward inclusion as part of the overall objective. Lastly, we hope to embed our approach for dynamic representation of virtual contents into existing adaptive MR placement systems [21]. While we focus on the specific representation of virtual content, such as combination would enable us to study higher-level actions such as transitions between environments or when no suitable host object is present.

View-dependent adaptation In our current implementation, we perform a single mapping to integrate the 2D application into the geometry of the physical target. This means the mapping does not adapt to users’ position in real-time. We believe that frequently changing this mapping would distract users. Our approaches can create ambient information displays, similar to decorative objects. Our approach allows users to get familiar with a specific setting, and to leverage their spatial memory. This goes in line with automatic approaches for placement of physical displays such as the approach by Fender et al. [37]. Future approaches might consider a more dynamic approach, where users’ position is continuously taken into account. We believe, however, that in our setting, infrequent changes are more beneficial for users since they are less distracting. Our approach is optimized for this setting and uses the geometry of the physical objects and the hierarchy of the 2D applications as main influencing factors, in contrast to prior approaches that largely focused on users’ viewing direction [38].

Legibility and clarity of MR content. Lastly, the output of our approach is limited by the resolution of the MR headset. While we use a high-resolution headset (Varjo XR-3), reading text and working with high-resolution graphics is still a major challenge for all MR headsets. The participants in our evaluation did not report challenges of legibility or clarity. Our algorithm is designed agnostic of display output, and can be adjusted to its capabilities by adjusting the minimum size constraint. We hope to test our approach with headsets that feature different resolutions in the future to explore the benefits and limitations of our system.

We believe that BlendMR presents a viable alternative and complement to current approaches to displaying MR interfaces. Our computational method to map 2D interfaces to physical objects provides a path to enabling users to take advantage of this representation by removing the need for manual design.

Chapter 4

Parametric Haptics: Versatile Geometry-based Tactile Feedback Devices

The following chapter is based on *Parametric Haptics: Versatile Geometry-based Tactile Feedback Devices*, submitted to *The ACM Symposium on User Interface Software and Technology 2023*.

4.1 Introduction

Haptic feedback is an important perceptual modality in our daily interactions with the world, but is difficult to reproduce artificially. The need for actuation makes careful mechanical design of the devices imperative. Many capable devices have been demonstrated that can, e.g., render shapes [39, 49, 10], textures [70], softness [115], impact [83], moving mass [112], and many more features. Such devices that provide active haptic feedback to the user are often presented in rigid form factors, such as hand-held controllers, body-grounded force feedback, wearables, and so on. These devices are usually specially designed for specific applications.




Figure 4-1: We present customizable and versatile haptic patches that are lightweight, reconfigurable, and can be integrated as wearables or into objects. (a) Users can customize their own haptic patches with a software aid we provide. (b) One linear pull can actuate many factors, producing haptic sensation on a large area of the skin. (c) Our approach can be integrated as wearables, into existing small objects, and designed for VR experiences.

The body is prime real estate for haptic feedback, thus wearable haptic feedback devices are a big area of interest in research. Ideal wearable haptic devices are lightweight, small, and conformable. Vibrotactile actuators have been used extensively due to their small form factor [60, 68, 79, 88]. To go beyond high-frequency stimuli of vibrotactiles, researchers have investigated other actuation techniques to produce indentation into the skin or stretch the skin laterally. Such tactile displays are driven by, e.g., linear resonant actuators [27], piezoelectric actuators [76, 86], custom magnetic actuators [102], shape memory alloys [47], or hydraulic electrostatic actuator arrays that form raised [75].

However, while the location where current tactile devices provide feedback at can be controlled dynamically, the *type of feedback* (e.g., vibration, indentation, skin stretch) cannot be adapted to new environments or scenarios. The feedback type is baked into the devices' design. We aim to overcome this by developing a device that allows for tactile feedback to be customizable and adaptable *after* it was designed and manufactured.

4.1.1 Parametric Haptics

In this paper, we present a modular haptic interface customizable to render different tactile stimuli. Figure 4-2 shows our device which consists of a (1) passive flexible and *versatile* tactile output geometry that can be connected to a (2) generic wearable actuation interface

(here, using micro gear motors).




Figure 4-2: (a) Our tactile feedback device consists of custom tactile patches that can be connected to a generic actuation interface, that (b) is wearable. (c) Multiple tactors are actuated to stroke over the skin, while the tactor design is user-defined.

The key benefit is that the tactors on our passive patches can be *custom designed* for new experiences and applications, e.g., for different virtual reality environments, assistive devices, or telepresence, while the generic actuation interface remains the same. Figure 4-2b shows how the patches are thin and flexible and can conform to users' bodies or objects. The custom tactors can be tailored to the specific application to render tactile feedback that is closely aligned with the expected experience. As shown in Figure 4-2c, multiple tactors are actuated by an integrated pulling geometry and can produce skin stretch, affective stroking, and indentation depending on the tactors' design. These patches are then 3D printed and connected to the actuation interface, e.g., the bracelet in Figure 4-2.

This decomposition into customizable tactile patches and generic actuation interface allows for tactile experiences to be shared by 3D printing as shared tactile patch geometry. For example, assets for VR environments could be delivered with their tactile experience in addition to their visual representation.

4.1.2 Contribution

The main contribution of this work is the design of versatile tactile feedback encoded in their geometry. The specific contributions of this paper are the following:

1. *Actuated geometry.* We develop a parametric geometry that is versatile and can be adapted to mimic custom tactile sensations.
2. *Tactile perception study.* We investigate the mapping between geometry parameters and how they are perceived. We find a good agreement, which supports the customizability of our passive tactile patches.
3. *Wearable prototypes.* Showcasing how the actuation geometry can be used in small wearable forms on varying locations with different sizes.
4. *Design space & applications.* We investigate the applicability of our parametric haptic device by showcasing distinct applications in VR, assistive technology, and notifications. We also outline the design space of our tactile patches.
5. *Design tool.* We provide a simple design tool that generates the actuation geometry for user-defined parameters.

4.2 Related Work

In this section, we incorporate literature on tactile perception as a basis for the use of 3D-printed geometric patches to convey varied tactile sensations. Further, we highlight the HCI contributions of passive, 3D printed mechanisms for delivering haptic feedback. Our work builds on creating customizable haptic interfaces that leverage geometric parameters that have promising applications to VR, accessible navigation, and tactile communication (e.g., notifications).

4.2.1 Tactile Perception

The human sense of touch is encoded through receptors in the skin known as mechanoreceptors. There are four distinct mechanoreceptors found in the hairless skin (i.e., palms, soles of feet). Pacinian Corpuscles sense high-frequency vibrations and Ruffini Endings detect skin stretch, Meissner Corpuscles sense the rate of skin deformation, and Merkel Cells detect spatial features. Meissner Corpuscles and Merkel Cells are not present in hairy skin [63, 46].

Hairy skin covers most of the skin’s surface and is capable of discriminative touch [93]. While it may have low spatial resolution, it is an effective site to convey both skin stretch and vibratory feedback. Prior work has demonstrated that skin stretch feedback can be used to provide directional cues [69, 8, 17, 66]. Vibratory feedback is salient for texture perception, namely perceived roughness [13, 12]. Surface geometry, which includes the density and shape of textural elements, also has importance for texture perception, namely roughness [73, 69].

The exact tactile cues that impact of perception of compliance (i.e., hardness and softness) are still unresolved [126]. When an object is indented into a stationary finger during passive touch (i.e., stimuli imparted on user by external agent instead of user motion), we are reliant on our tactile, or cutaneous, cues [48, 111, 117, 116]. Recent work by D’Aurizio et al. uses an on-forearm, wearable haptic device for skin-stroking and squeezing. This work demonstrates that users are able to discriminate different compliances successfully [30] and showcases the promise for the use of passive tactile feedback to enable the discrimination of relative hardness and softness.

Touch can be categorized as discriminative, which we described earlier, as well as affective [93, 3]. Recent findings uncovered an additional mechanoreceptor, C-tactile afferents (CTs), which are only present in the hairy skin. CTs are selectively responsive to stroking [93]. Research by Hertenstein demonstrates that emotional communication (i.e., love, anger, sadness) can be conveyed through touch. [52]. Furthermore, people are able to differentiate between different emotions through touch alone [52].

The literature suggests that the hairy skin is a prime location to impart a varied and distinct set of tactile sensations. We present the design of 3D-printed patches that are versatile in the type and density of geometric parameters and can be actuated over a large area to elicit a variety of tactile sensations that can support a range of applications.

4.2.2 Portable and Wearable Haptic Feedback Devices

A significant amount of prior work exists on haptic feedback devices that are portable enough to be worn or carried [14]. Handheld devices such as [121, 11, 23] are controllers that complement VR scenarios with haptic feedback for touching or grasping virtual surfaces, textures, and shapes. Handheld devices, however, limit haptic feedback to the hand and do not allow free-hand interactions, making them less applicable to use outside of VR.

Wearable haptic devices can be mounted on the body, freeing a user’s hands. Such devices could be used to render both force and tactile feedback. Force feedback of interacting with virtual objects can be rendered with an exoskeleton-device arresting forces on users’ various body parts [36, 4, 53, 22]. These devices rely on the body part’s natural degree of freedom and could sometime limit the freedom of movement, therefore needing much design consideration.

Vibrotactile is commonly used in wearable haptic devices to render tactile feedback due to the small form factor and cheap price of vibratory actuators. Researchers have used vibrotactile wearables to render pleasant social touch [60], provide navigation cues [68, 79], convey textual information [88], and so on. Electrotactile allows wearable haptic devices to be miniaturized to as small as a temporary tattoo [122]. Despite their small form factors, vibrotactile and electrotactile stimuli remain limited: electrotactile stimuli can not easily stimulate Pacinian Corpuscles (PC) laying in deeper region of the skin [65], and could give an unnatural ”electric feeling” [71], while PC dominates vibrotactile perception and requires vibrotactile stimuli to be placed far apart due to its large receptive field [8].

On the other hand, physiology research has shown that skin strain allows mechanoreceptors to respond more quickly and accurately [31, 32, 100]. Leroy et al. proposed using

flexible hydraulically amplified electrostatic actuators to produce normal and shear forces to the user’s body [75]. Teng et al. suggested providing haptic feedback to mixed reality environments with a device that presses against the user’s fingerpad when they touch something;

Wearable tactile feedback devices can be used to provide learned tactile vocabularies [64] (e.g., notifications) or to provide directional cues [69, 8, 17, 66]. Research by Bark et al. suggests that shear force when delivered to the arm provides superior directional cues compared to that of vibrotactile [8]. Furthermore, skin stretch can also be used to create convincing illusions, such as stiffness perception [104].

4.2.3 Embedding Functionality Within Geometry

Passive haptic feedback mechanisms dissipate applied mechanical energy through interaction with a human [128]. Recent work has demonstrated the significance of embedding functionality within passive mechanisms to provide haptic feedback [128, 78]. Chang et al. introduced simple-to-fabricate Kirigami structures with embedded geometric patterns used to render different types of haptic feedback [18]. Complementary to active feedback mechanisms, researchers also introduced passive haptic feedback devices, e.g., hand-held force feedback controlled by compliant mechanisms [78, 128], or tactile feedback encoded in 3D printed textures [118, 55, 54]. The absence of actuation makes passive haptic feedback devices less complex to build at the expense of dynamic control.

4.3 Designing Parametric Haptics

We aim to develop a new form of haptic wearable that is highly-customizable, lightweight, and requires little electronics. We leverage 3D printing technology to fabricate such wearable patches. A small pull force from a motor can simultaneously actuate many tactors at once, allowing them to produce skin stretch and stroking sensations on the users’ skin, rendering various complex tactile experiences. The simplicity of our actuation geometry further allows

the wearable patches to be customized for different applications, be worn on different body parts (e.g., arms, legs, torso), or integrated with existing objects (e.g., glasses, wrapped around water bottles, desktop surface).

In this section, we show the design of our tactile patches. Figure 4-3 illustrates the main components that make up our tactile patch. The custom-designed tactors touch the user’s skin when actuated. The tactors are connected with an actuation geometry that we detail in the following. The pull string moves all connected tactors approximately perpendicular, such

4.3.1 Actuation Geometry

The actuation in our tactile patches is inspired by tendon-driven mechanisms. Instead of assembling threads, which is tedious, we develop a geometry that replaces the literal string with printable structures and can be entirely 3D printed.

We design stem-like structures, shown in Figure 4-3, connect one tactor to the next, and extend out to act as a pull string after the last tactor. These stems propagate the external pulling force to actuate all tactors that are connected to them.

To allow the tactors to rotate, we anchor their position on the patch through compliant hinges, i.e, thin struts that flex easily (0.6 mm). We also place compliant hinges between the tactors and the stem. The hinge at the front of the tactor is a simple thin flexure. To connect to the next tactor, the stem structure extends from behind the tactor with a short thin curved section. When the structure is actuated, this curved section straightens out to be orthogonal to its adjacent tactor. This hinge placement transforms the linear pulling motion into a rotation of the tactor.

In our experiments, placing the stem hinge min. $1/2$ its length on a given tactor design yields reliable results. We use a clearance of 0.2 mm on both sides between the tactor and the stem to avoid friction and smooth actuation. We found that making the stem structure’s width around $1/3$ the width of the tactors, and its thickness around half its width results in




Figure 4-3: (a) An unactuated haptic patch with geometry recommendations for each part. (b) The same patch with its tactors actuated.

reliable prints yet keeps the required pulling force low. These recommendations were tested on tactor sizes of 5-20 mm wide and for up to 16 connected tactors. These parameters might have to be adjusted for different tactor geometries (e.g., thicker tactors may require a stronger stem structure to actuate).

4.3.2 Fabrication

The tactile patches are printed with thermoplastic polyurethane (TPU), a common rubber-like filament. We use the filament *NinjaFlex* (Shore 85A). We choose TPU due to its flexibility allowing patches to easily conform to the body (e.g. wear on forearm), or daily objects (e.g. wrap around a water bottle). Due to the complexity of our tactile geometry, we need support material, for which we use Polyvinyl alcohol (PVA by Polymaker), which is an off-the-shelf water-soluble support material. We use a dual-extruder commercial printer (Tenlog TL-D3 Pro) to print the patches.

Our TPU can stretch substantially (up to 660%) which is advantageous for the conformability of our patches, but not for propagating the pulling force through the actuation geometry. In early prototypes, energy was lost in the stretch of the material, meaning that tactors were actuated less and less. To mitigate that, we generate a thin line (0.4 mm width) embedded into the actuation geometry which we print with our rigid support material. Because the PVA is fully enclosed within the TPU, it does not dissolve and provides rigidity to our actuation geometry.

4.3.3 Design Tool




Figure 4-4: Users can input parameters of the haptic patch they want to create, an design tool automatically generates a working actuation geometry, along with a correctly angled guidance plane for users to design their own tactors.

In general, tactors are relatively flat geometries, allowing them to lay down and not contact the skin when they are not actuated. It is to note that even when tactors are not actuated, they are not laying absolutely flat, but are slightly angled upward along with the front stem structure it is connected to. Such angles allow all tactors to be placed on the same horizontal line, i.e. none is higher or lower than another. This angle is calculated with the aiding software component, which helps users automatically generate working actuation geometries, without having to do much deliberate calculations (e.g. for the aforementioned curvature). Our software is a plugin for Rhino 6, with the backend written in Python 3.6 and Rhino.Python API, and front-end built with Eto.Forms.

The software takes user-defined parameters as input and generates a working actuation

geometry, including stem structures, stem hinges, anchored tactor-rotation hinges, return spring, the bottom patch, and drawing planes at a calculated angle for users to design their own tactors on. Users can easily input desired parameters and leverage the generated geometry to design their own haptic patches for various use scenarios.

4.3.4 Motorized Actuation Interface

The wearable design interface is a modular design composed of horizontal motor mounts to accommodate a single motor. In Figure 4-5, we show the components of our wearable bracelet. Motors are spaced 13 millimeters apart (the width of the motor), and placed side-by-side to maximize the number of actuators per surface area. Wearable patches are attached to rails that mate to complementary slots at the base of each motor mount. This modular design allows for seamless exchange of patches of various sizes and tactor densities, as well as to easily adapt the form factor for placement on different areas of the body (i.e., forearm, back, etc).

We use 12 mm DC gear motors (Ruimou 6V 100 RPM), which are controlled by an Arduino Nano microcontroller through DRV8833 motor drivers. Wireless versions are attainable with little engineering effort by embedding WiFi-enabled microcontrollers and batteries (e.g., Seeed Studio Xiao and 7.4V LiPo battery as illustrated on the left of Figure 4-5). Each motor features a pulley system. We use fishing line (Power Pro Spectra Fiber), which is fed through a guidance hole at the base of the motor mount and attaches to a hook on the patch pull string.

For our demonstration, we implement a simple open-loop control strategy. We linearly actuate the motors by varying the direction of the rotation, the Pulse-Width-Modulation (PWM), and the duration and frequency of the pulses, allowing us to approximate control over the displacement of the tactors and velocity of actuation to achieve a diverse set of tactile sensations. Control strategies can be adapted to elicit desired tactile feedback and incorporate refined control with appropriate hardware considerations (e.g., encoder for refined position or velocity control, etc).




Figure 4-5: Modular wearable interface. Users can exchange patches of different sizes for various use scenarios with the same wearable interface.

4.4 Design Space

Our tactile geometry spans a larger design space, which we show in Figure 4-6. We structure it into the following dimensions.




Figure 4-6: Our haptics patches have a wide design space for creating different tactile sensations.

Tactor geometry. Tactors are what contacts the skin and create skin stretch and stroking when actuated. They can take nearly arbitrary forms to render various sensations. We show only 4 examples of different tactor shapes in Figure 4-6, but designers can create arbitrary tactors. Tactor parameters. The tactor geometry can be further parameterized. To keep the overall sensation, but alter, e.g., its intensity, designers can change the tactors size,

particularly their length, and thickness. Thicker tactors will render a harder stimulus, longer tactors may render softer, more stroke-like sensations.

Tactor parameters. The tactor geometry can be further parameterized. To keep the overall sensation, but alter, e.g., its intensity, designers can change the tactors size, particularly their length and thickness. Thicker tactors will render a harder stimulus, longer tactors may render softer, more stroke-like sensations.

Density. Our tactile patches are designed to render feedback on larger areas of the skin. To do so, multiple tactors and tactor rows can be combined within one patch. The designer may vary their density to render, e.g., individual stimuli with sparse tactors or area sensations using dense tactor layouts. The density can be varied non-uniformly as well.

Heterogeneity. Different tactor geometries may be combined within one patch, either row-wise or entirely free.

Direction. The direction in which the tactors are actuated can be designed as well. All tactors might be actuated in the same direction. Designers might actuate tactor rows in opposite directions by using two actuation interface bracelets. The tactors may also be arranged in curved paths on the patches to vary the direction of the stimuli within the patch.

4.5 User Study: Perception of Geometric Parameters

We investigate whether the geometry parameters of our approach match the corresponding perceptual sensation. We test three different parameters (*tactor*, *stiffness*, *layout*), with a total of 18 differently parameterized tactile patches. Participants experienced the individual patches, and rated them for different sensations such as softness, coarseness, or sharpness. Results indicate a good agreement between our parameters and the reported sensations.

4.5.1 Study Design

We use a within-subjects design with three independent variables for different parameterizations of the tactile patches: tactor with three levels (triangle, round, bristles), stiffness with three levels (soft, medium, hard), and layout with four levels (sparse, medium, dense, gradient).

Design of tactile patches

The independent variables are reflected in the design of the tactile patches, all shown in Figure 4-7. The tactors and layout are chosen to cover a range of possible sensations. The parameterization of stiffness for each tactor varies. For triangle and round, stiffness is controlled through the thickness of the tactor. For the bristles, stiffness is controlled by the length of the bristles. The layout is controlled by varying the spacing between individual tips over the length of each patch (total length 70 mm; sparse: 4 tips, medium: 8 tips, dense: 16 tips, gradient: 8 tips with the distance between tips increasing by 0.7 mm each tip).

Data collection

We collected data on participants' tactile sensations. After experiencing each patch they were asked to rate the sensation on different seven-point scales from soft to hard, over-the-skin to into-the-skin, blunt to sharp, and single contacts to many contacts. We further elicited qualitative feedback using semi-structured interviews.

4.5.2 Participants and Apparatus

We recruited 12 paid participants (6 male, 6 female), aged $M = 25.92$ years ($SD = 2.50$), all students and staff from a local University through convenience sampling. We compensated them with \$15 for their participation. All participants felt the patches with their dominant




Figure 4-7: Independent variables are reflected in tactors and their layouts. We tested 18 patches in total.

arm, except for one participant wearing braces on their dominant arm. The study was conducted in an experimental space, shown in Figure 4-8. Each tactile patch was embedded into a laser-cut box. The distance between participants' arms and the patches could be adjusted with modular spacers to ensure proper contact, described below. All patches were mounted onto a wooden table, and covered with a visual barrier so that participants could not see the individual tactors.




Figure 4-8: Study setup. Before conducting the study, each participant’s arm is first calibrated for consistent indent of tactors into the skin.

4.5.3 Procedure

After completing the consent form and background questionnaire, participants were briefed on the purpose of the study. We then performed a height calibration. Participants were asked to place their dominant arm onto a transparent box, shown in Figure 4-8. The experimenter selected a 3D-printed spacer so that the tactors of the test patch indented participants’ skin between 1 - 2 mm. The spacer was then used for all other patches throughout the experiment.

The presentation order of the 16 patches was randomized. For each patch, the experimenter triggered the tactile sensation multiple times by pulling a string at a constant velocity to a distance of 1.5 mm. Participants were then asked to rate the experience on aforementioned scales, and briefly describe the sensation.

After experiencing all patches, participants were asked to reflect on the sensations and provide examples of what the experience reminded them of, either in real life or digital experiences such as games or VR. The experiment took approximately 45 minutes per participant.

4.6 Results

In summary, subjective ratings indicate that triangle tactors were perceived as harder and more sharp, and as more protruding into the skin. Bristles and round tactors led to a softer, surface-oriented sensation. The parameterized stiffness of the tactors successfully changed participants' perception towards soft, blunt and unified sensations. Dense layouts move participants' reported perception into a similar direction. This is also reflected in participants' qualitative comments.




Figure 4-9: Mean subjective ratings from the experimental questionnaires. Error bars indicate standard errors, (*) indicates statistical significant differences.

4.6.1 Quantitative Rating of Stimuli

We analyzed the subjective ratings on the perception of soft to hard, over skin to into skin, blunt to sharp, and single to many contacts, all on scales from 1 to 7 (e.g., one scale ranged from 1 - soft to 7 - hard). Results are shown in Figure 4-9.

We performed an Aligned Rank Transform (ART) [123] on the ranking using the Windows tool of ARTool. The transformed data was analyzed using a series of repeated measures ANOVAs, performed in SPSS 29 . Since the data was unbalanced between the three independent variables, we performed individual RM-ANOVAs on tactor and stiffness; and tactor and layout. For clarity of exposure, we only report on significant main effects ($\alpha = 0.05$). Posthoc tests of significant main effects were performed including Bonferroni adjustment for multiple comparisons.

For *soft to hard ratings*, we found significant main effects for tactor ($F(2,22) = 11.323$, $p \leq 0.001$), stiffness ($F(2,22) = 15.727$, $p \leq 0.001$) and layout ($F(3,33) = 9.796$, $p \leq 0.001$). Bristle and round tactors were perceived as significantly softer than triangle tactors (all $p \leq 0.05$). Soft tactors (i.e., tactors with thin tips or long bristles) were perceived as significantly softer than medium or hard tactors (all $p \leq 0.01$), and dense layouts were perceived as softer than medium and sparse layouts (all $p = 0.001$).

For *over-skin to into-skin ratings*, we found significant main effects for tactor ($F(2,22) = 12.371$, $p \leq 0.001$), stiffness ($F(2,22) = 7.763$, $p = 0.003$), and layout ($F(3,33) = 4.558$, $p = 0.009$). Triangle tactors were perceived as delivering more of an into-skin sensation compared to bristles ($p = 0.007$) and round tactors ($p = 0.010$), similar to hard tactors compared to soft ones ($p = 0.008$). Dense layouts were perceived more as over-skin sensations compared to medium and sparse layouts (all $p \leq 0.05$).

For *blunt to sharp ratings*, we found significant main effects for tactor ($F(2,22) = 36.343$, $p \leq 0.001$), stiffness ($F(2,22) = 14.762$, $p \leq 0.001$) and layout ($F(3,33) = 10.354$, $p \leq 0.001$). Triangle tactors were perceived as sharper compared to bristles and round tactors (both $p \leq 0.001$). Soft tactors were perceived as more blunt compared to medium ($p = 0.001$) or

hard tactors ($p = 0.005$). Dense layouts were perceived as more blunt sensations compared to medium and sparse layouts (all $p \leq 0.01$).

For *single vs multiple sensations ratings*, results indicate a main effect for stiffness ($F(2,22) = 10.412$, $p \leq 0.001$), but no main effect for tactors ($p = 0.924$) or layout ($p = 0.144$). Soft tactors were perceived to yield sensations closer to a single surface compared to medium ($p = 0.004$) or hard ($p = 0.016$).

4.6.2 Qualitative feedback

In general, participants' qualitative feedback indicates that our tactile patches can render a wide range of sensations for various usage scenarios.

Influence of tactor

Triangle tactors were connected to scratching or stroking from people, animals, or hard objects. *"Like finger nails scratching"* (P1 - sparse triangle); *"dog paw or cat paw scratching"* (P8 - sparse triangle). Stiff triangle tactors were related to unpleasant scratching that could be used for urgent notifications: *"...pine needles...that hurts"* (P2 - hard triangle); *"try to alert the user... or bring back their attention..."* (P6 - medium triangle). Participants related coarse and soft triangle tactors to more pleasant sensations ($n = 5$) such as a comfortable scratch of the itch ($n = 3$) and a feeling in-between scratch and touch from a human or a pet ($n = 2$).

Round tactors were associated with scratching or stroking, but pleasant and blunt: *"between finger and nails... between touch and scratch"* (P2 - medium round); *"more like a brush in a narrow area"* (P10 - medium round); *"...a VR scenario where I'm a pet being groomed"* (P12 - sparse round). Further comments suggest that round tactors can render versatile sensations, such as big velcro (P1 - sparse round), rub against wood surface (p6 - hard round), water wave (P7 - hard round), or gentle grabbing (P2 - soft round).

Bristle tactors produced the widest range of haptic sensations, ranging from neutral to pleasant. Participants compared the sensations with social and affectionate touch ($n = 3$), and with different bristle structures such as a comb (P18), toothbrush (P9, P11), and shoe brush (P8). Bristle tactors produced the sensations of friction: *“it’s like when you have a light comfortable wearable, and when you are like running, it rubs against your skin”* (P6 - short bristles); Long bristles triggered sensations such as *“a soft object hitting you fast”* (P6); a bubble wrap surface (P11), and a blunt tapping sensation from an eraser (P12).

Influence of stiffness

Softer tactors and sparse layouts were suggested for subtle notifications ($n = 5$), haptic feedback for button press (P7 - sparse triangles), replacement for vibration (P7, P12), or augmenting digital media with pleasant sensations ($n = 3$). Harder patches were suggested for rendering unpleasant sensations such as cuts ($n = 3$).

Influence of layout

Gradient layouts were often associated with motion: *“Yeah I can feel a single thing that is moving, it’s kind of like a point”* (P3 - gradient triangle); *“soft fine sand... sliding off”* (P2 - gradient bristles); *“Wind. Like spring wind. When I sleep, when I was young, I was sleeping during the spring and then I open the window.”* (P7 -gradient bristles); *“running something along your arm”* (P12 - gradient round). Dense layouts decreased the perception of motion, and participants’ comments were focused on the texture or material qualities: *“it’s like a shower loofah”* (P1 - dense bristles); *“someone’s dry skin... or maybe plywood”* (P2 - dense round); *“like, you know how you feel when it’s dry and your skin is scratched off”* (P6 - dense triangle); *“a pool of tiny stones”* (P7 - dense triangles); *“needles in a square or a circle”* (P10 - dense triangles).

4.7 Application Examples

We showcase the versatility of our haptic patches with several application scenarios, including haptic feedback for VR, integration with existing objects to communicate information through haptics, and a wearable notification device that can send users different kinds of notifications.




Figure 4-10: (a) A haptic patch for a VR jungle expedition game (b) a user being swept by leaves in the VR game, feeling the sensation on their forearm (c) a user being bitten by spiders in the VR game, with the haptic patch rendering this sensation.

4.7.1 Haptic feedback in Virtual Reality

With our modular design, haptic sensations that accompany MR environments can be shared as assets. Creators can upload on Thingiverse, for example, and users can download and print them.

Here we showcase the example of a haptic patch created for a jungle expedition game (Figure 4-10). The patch includes three kinds of tactors: a stiff triangular laid out densely to render a sharp cut, a dual-triangular tactor with a thin connector in the middle that pinches the user when actuated, and a soft feather-like tactor to render stroking from relatively soft objects.

During the VR jungle expedition game, the user is taking a walk through the jungle. Their arm is swept against a bush by the side, the stroking tactors are actuated at this time to produce the corresponding haptic feedback. Then, out of the spirit of exploration, the user places their arm into a hole. Spiders climb out of the hole and onto the user’s arm, biting their skin. The pinching tactors are now actuated to give this encounter more realism. Afterward, the user accidentally gets a thin cut from a sharp plant. The stiff triangular tactors are actuated to render this cut.

Despite this example only showing one scenario for each kind of tactor’s usage, tactors could be reused for multiple scenarios. Multi-modal perception studies indicate that the visual sense is much stronger among different kinds of perception [74, 44, 120]. This means that individual tactors can be used to render multiple related VR encounters. This goes in line with the findings of our study, where participants associated different sensations with the same tactor.

4.7.2 Augmenting objects

Parametric Haptic interfaces are small and flexible, and can be integrated with existing objects. In Figure 4-11 we show a simple 3-tactor patch that is integrated into a pair of glasses, without having to compensate for the form factor. If connected with smart

devices or sensing systems, this small haptic interface could quickly and effectively deliver notifications or low-level information to the user, while being unobtrusive.



Figure 4-11: (a) A row of 3 tactors integrated into an existing pair of glasses (b) a user wearing a pair of Parametric Haptic-integrated glasses in daily life, the glasses can unobtrusively and effectively deliver information.

4.7.3 Wearable notification device

Parametric Haptic patches could be further designed into various wearable devices. Here we showcase a wearable that is worn on the user’s back that can render different kinds of tactile notifications. Tactors in this example patch are relatively larger and placed further apart, taking advantage of the large area of the back and the fact that the mechanoreceptors on the back have larger receptive fields and are of low density [7, 69]. The patch shown in Figure 4-12 has four curved rows of round tactors of medium stiffness to deliver navigation instructions [64]. The ability of our actuation geometry to redirect tactors allows motors to be placed only on two sides of the haptic patch, while allowing the patch to render four directions to help users navigate. Such usage could be leveraged for accessibility purposes (e.g. aiding low-vision users with navigation).

Besides directional information, this haptic patch could also deliver users various notifications. Informed by our study and previous works [69, 8, 17, 66], a stroking sensation from a soft tactor is perceived as affective and comforting [93, 52]. Here, a row of soft bristle tactors are placed on the left side of the patch (Figure 4-12). They can be actuated when the user

receives a text notification from friends and family - one of many scenarios that could be suitable for affective touch. For a more urgent notification such as a phone call from users' social circle, a shorter row of elliptical tactors are used to render a poking sensation (insight from participants in our study). Another short row with softer tactors could be used for other subtle notifications, potentially replacing vibration notifications. The row of triangular tactors on the right side of the patch (Figure 4-12), on the other hand, would render more urgent notifications that need user's immediate attention (e.g. fraudulent charge on credit card).




Figure 4-12: (a) This back wearable patch can deliver various notifications and information to the user's lower back. (b) The user gets a gentle notification for a text message from their family. (c) We show the various tactors and directions we can render on this larger area.

4.8 Discussion

We presented Parametric Haptics, a versatile and customizable tactile feedback device. The versatility of the devices stems from its decomposition into (1) passive 3D printed tactile feedback patches with custom tactors designs and (2) a generic actuation platform that the tactile patches can be connected to. Our design has several tangible benefits, i.e., that the 3D printed customizable geometry allows them to be designed for versatile applications. Since the patches are thin and soft, they conform easily to the user’s body or objects. Additionally, the tactile patches are designed to provide tactile feedback over a larger area while offsetting the actuation away from the area of feedback.

Conceptually, the biggest benefit is that the tactors can be designed to match the represented stimuli closely and are easily exchangeable. Since the stimuli are encoded in the geometry, this allows the tactile experience to be shared with and replicated by others—a longstanding challenge in the field of haptics.

The current scope of this work presents some limitations. The distance between the tactor and the skin will influence the perception. While we have maintained a consistent distance in our user study, we have not investigated different distances and their influence. Our user study shows differences in the perception of lateral skin stretch and indentation normal to the skin, and we believe that we further enhance this and produce more distinct differences by, e.g., building a Sarrus linkage¹ as a tactor mechanism to produce true indentation. We also believe that we can miniaturize the tactors by using high-resolution fabrication techniques (e.g., stereolithography 3D printing) and investigate the use of these patches on other body parts with higher density in mechanoreceptors. In the future, we want to dig deeper into the design space that this paper opened and into the aforementioned parameters and their influence on perception.

¹https://en.wikipedia.org/wiki/Sarrus_linkage

Chapter 5

Discussion and Conclusion

With the two works presented in this thesis, I showcased how bidirectional interactions could be created between the digital and physical worlds: BlendMR brought the unlimited access to information from the digital domain to the physical, while ParametricHaptics brought rich sensations from the physical world to the digital. Many opportunities exist in bringing more advantages of one domain into another, e.g. how could we gain simple and remote control of existing everyday physical objects as how we could control digital elements on an interface? Moreover, one could also explore what are meaningful ways to connect the two worlds, e.g. leveraging semantic associations [21], or even utilizing seemingly disadvantages of one world along with imperfections of human perception to create interactions non-existing in the other [1, 2].

Bibliography

- [1] Parastoo Abtahi and Sean Follmer. Visuo-haptic illusions for improving the perceived performance of shape displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.
- [2] Parastoo Abtahi, Mar Gonzalez-Franco, Eyal Ofek, and Anthony Steed. I’m a giant: Walking in large virtual environments at high speed gains. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [3] Rochelle Ackerley, Karin Saar, Francis McGlone, and Helena Backlund Wasling. Quantifying the sensory and emotional perception of touch: differences between glabrous and hairy skin. *Frontiers in Behavioral Neuroscience*, 8, 2014.
- [4] Ahmed Al Maimani and Anne Roudaut. Frozen suit: Designing a changeable stiffness suit and its application to haptic games. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI ’17, page 2440–2448, New York, NY, USA, 2017. Association for Computing Machinery.
- [5] Abul Al Arabi, Jiahao Li, Xiang ’Anthony Chen, and Jeeun Kim. Mobiot: Augmenting everyday objects into moving iot devices using 3d printed attachments generated by demonstration. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI ’22, New York, NY, USA, 2022. Association for Computing Machinery.
- [6] Ronald T. Azuma and Chris Furmanski. Evaluating label placement for augmented reality view management. In *2003 IEEE / ACM International Symposium on Mixed*

- and *Augmented Reality (ISMAR 2003)*, 7-10 October 2003, Tokyo, Japan, pages 66–75. IEEE Computer Society, 2003.
- [7] Paul Bach-y Rita. Tactile sensory substitution studies. *Annals of the New York Academy of Sciences*, 1013(1):83–91, 2004.
 - [8] Karlin Bark, Jason W Wheeler, Sunthar Premakumar, and Mark R Cutkosky. Comparison of skin stretch and vibrotactile stimulation for feedback of proprioceptive information. In *2008 Symposium on haptic interfaces for virtual environment and teleoperator systems*, pages 71–78. IEEE, 2008.
 - [9] Blaine Bell, Steven Feiner, and Tobias Höllerer. View management for virtual and augmented reality. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology, UIST 2001, Disney’s BoardWalk Inn Resort, Walt Disney World, Orlando, Florida, USA, November 11-14, 2001*, pages 101–110, New York, NY, 2001. ACM.
 - [10] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. Normaltouch and texturetouch: High-fidelity 3d haptic shape rendering on handheld virtual reality controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology, UIST ’16*, page 717–728, New York, NY, USA, 2016. Association for Computing Machinery.
 - [11] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. Normaltouch and texturetouch: High-fidelity 3d haptic shape rendering on handheld virtual reality controllers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology, UIST ’16*, page 717–728, New York, NY, USA, 2016. Association for Computing Machinery.
 - [12] Sliman Bensmaïa and Mark Hollins. Pacinian representations of fine surface texture. *Perception & psychophysics*, 67:842–854, 2005.
 - [13] Sliman J Bensmaïa and Mark Hollins. The vibrations of texture. *Somatosensory & motor research*, 20(1):33–43, 2003.

- [14] Carlos Bermejo and Pan Hui. A survey on haptic technologies for mobile augmented reality. *ACM Comput. Surv.*, 54(9), oct 2021.
- [15] Mark Billinghurst, Hirokazu Kato, Ivan Poupyrev, et al. Tangible augmented reality. *ACM Siggraph Asia*, 7(2):1–10, 2008.
- [16] John Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
- [17] Nathaniel A Caswell, Ryan T Yardley, Markus N Montandon, and William R Provancher. Design of a forearm-mounted directional skin stretch device. In *2012 IEEE Haptics Symposium (HAPTICS)*, pages 365–370. IEEE, 2012.
- [18] Zekun Chang, Tung D Ta, Koya Narumi, Heeju Kim, Fuminori Okuya, Dongchi Li, Kunihiro Kato, Jie Qi, Yoshinobu Miyamoto, Kazuya Saito, et al. Kirigami haptic swatches: Design methods for cut-and-fold haptic feedback mechanisms. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020.
- [19] Jieshan Chen, Amanda Swearngin, Jason Wu, Titus Barik, Jeffrey Nichols, and Xiaoyi Zhang. Towards complete icon labeling in mobile applications. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI ’22, New York, NY, USA, 2022. Association for Computing Machinery.
- [20] Yi Fei Cheng, Tiffany Luong, Andreas Fender, Paul Strel, and Christian Holz. Comfortable user interfaces: Surfaces reduce input error, time, and exertion for tabletop and mid-air user interfaces. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, ISMAR ’22, 2022.
- [21] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. Semanti-cAdapt: Optimization-based adaptation of mixed reality layouts leveraging Virtual-Physical semantic connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST ’21, pages 282–297, New York, NY, USA, October 2021. Association for Computing Machinery.

- [22] Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. Gravity: A wearable haptic interface for simulating weight and grasping in virtual reality. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, page 119–130, New York, NY, USA, 2017. Association for Computing Machinery.
- [23] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. Claw: A multifunctional handheld haptic controller for grasping, touching, and triggering in virtual reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–13, New York, NY, USA, 2018. Association for Computing Machinery.
- [24] Shakiba Davari, Feiyu Lu, and Doug A. Bowman. Validating the benefits of glanceable and context-aware augmented reality for everyday information access tasks. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 436–444, 2022.
- [25] Niraj Ramesh Dayama, Kashyap Todi, Taru Saarelainen, and Antti Oulasvirta. GRIDS: Interactive layout design with integer programming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pages 1–13, New York, NY, USA, April 2020. Association for Computing Machinery.
- [26] Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hirschman, Daniel Afargan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. Rico: A mobile app dataset for building Data-Driven design applications. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, pages 845–854, New York, NY, USA, October 2017. Association for Computing Machinery.
- [27] Artem Dementyev, Alex Olwal, and Richard F. Lyon. Haptics with input: Back-emf in linear resonant actuators to enable touch, pressure and environmental awareness. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, UIST '20, page 420–429, New York, NY, USA, 2020. Association for Computing Machinery.

- [28] Stephen DiVerdi, Tobias Hollerer, and Richard Schreyer. Level of detail interfaces. In *Proceedings of the 3rd IEEE/ACM International Symposium on Mixed and Augmented Reality*, ISMAR '04, pages 300–301, USA, November 2004. IEEE Computer Society.
- [29] Ruofei Du, Alex Olwal, Mathieu Le Goc, Shengzhi Wu, Danhang Tang, Yinda Zhang, Jun Zhang, David Joseph Tan, Federico Tombari, and David Kim. Opportunistic interfaces for augmented reality: Transforming everyday objects into tangible 6dof interfaces using ad hoc ui. *CHI EA '22*, New York, NY, USA, 2022. Association for Computing Machinery.
- [30] N. D'Aurizio, T. Lisini Baldi, A. Villani, K. Minamizawa, Y. Tanaka, and Domenico Prattichizzo. Wearable haptics for object compliance discrimination through passive touch. In *2021 IEEE World Haptics Conference (WHC)*, pages 894–899, 2021.
- [31] Benoni B Edin. Quantitative analyses of dynamic strain sensitivity in human skin mechanoreceptors. *Journal of neurophysiology*, 92(6):3233–3243, 2004.
- [32] Benoni B Edin and Niclas Johansson. Skin strain patterns provide kinaesthetic information to the human central nervous system. *The Journal of physiology*, 487(1):243–251, 1995.
- [33] Barrett Ens, Benjamin Bach, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, Arnaud Prouzeau, Christoph Anthes, Wolfgang Büschel, Cody Dunne, Tim Dwyer, Jens Grubert, Jason H. Haga, Nurit Kirshenbaum, Dylan Kobayashi, Tica Lin, Monsurat Olaosebikan, Fabian Pointecker, David Saffo, Nazmus Saquib, Dieter Schmalstieg, Danielle Albers Szafir, Matt Whitlock, and Yalong Yang. Grand challenges in immersive analytics. In *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*, pages 459:1–459:17, New York, NY, 2021. ACM.
- [34] Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. Spatial constancy of Surface-Embedded layouts across multiple environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, SUI '15, pages 65–68, New York, NY, USA, August 2015. Association for Computing Machinery.

- [35] Barrett Ens, Eyal Ofek, Neil D. B. Bruce, and Pourang Irani. Spatial constancy of surface-embedded layouts across multiple environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction, SUI 2015, Los Angeles, CA, USA, August 8-9, 2015*, pages 65–68, New York, NY, 2015. ACM.
- [36] Cathy Fang, Yang Zhang, Matthew Dworman, and Chris Harrison. Wireality: Enabling complex tangible geometries in virtual reality with worn multi-string haptics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–10, New York, NY, USA, 2020. Association for Computing Machinery.
- [37] Andreas Rene Fender, David Lindlbauer, Philipp Herholz, Marc Alexa, and Jörg Müller. HeatSpace: Automatic placement of displays by empirical analysis of user behavior. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, pages 611–621, New York, NY, USA, 2017. Association for Computing Machinery.
- [38] Andreas Rene Fender and Jörg Müller. SpaceState: Ad-Hoc definition and recognition of hierarchical room states for smart environments. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS '19, pages 303–314, New York, NY, USA, November 2019. Association for Computing Machinery.
- [39] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. in-form: dynamic physical affordances and constraints through shape and object actuation. In *Uist*, volume 13, pages 2501–988. Citeseer, 2013.
- [40] Krzysztof Z Gajos, Daniel S Weld, and Jacob O Wobbrock. Automatically generating personalized user interfaces with supple, 2010.
- [41] Ran Gal, Lior Shapira, Eyal Ofek, and Pushmeet Kohli. FLARE: Fast layout for augmented reality applications. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 207–212. ieeexplore.ieee.org, September 2014.
- [42] Christoph Gebhardt, Brian Hecox, Bas van Opheusden, Daniel Wigdor, James Hillis, Otmar Hilliges, and Hrvoje Benko. Learning cooperative personalized policies from

- gaze data. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, pages 197–208, New York, NY, USA, October 2019. Association for Computing Machinery.
- [43] Google/Ipsos, U.S. How People Discover, Use, and Stay Engaged With Apps, 2016. Retrieved Sept. 14 2022.
 - [44] Steve Guest and Charles Spence. Tactile dominance in speeded discrimination of textures. *Experimental brain research*, 150:201–207, 2003.
 - [45] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2022.
 - [46] K.S. Hale and K.M. Stanney. Deriving haptic design guidelines from human physiological, psychophysical, and neurological foundations. *IEEE Computer Graphics and Applications*, 24(2):33–39, 2004.
 - [47] Nur Al-huda Hamdan, Adrian Wagner, Simon Voelker, Jan Borchers, and Jürgen Steimle. Demonstration of springlets: Expressive, flexible and silent on-skin tactile interfaces. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–4, 2019.
 - [48] Roland Harper and SS Stevens. Subjective hardness of compliant materials. *Quarterly Journal of Experimental Psychology*, 16(3):204–215, 1964.
 - [49] Chris Harrison and Scott E. Hudson. Texture displays: A passive approach to tactile presentation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, page 2261–2264, New York, NY, USA, 2009. Association for Computing Machinery.
 - [50] Paul S Heckbert. Survey of texture mapping. *IEEE computer graphics and applications*, 6(11):56–67, November 1986.
 - [51] Steven Henderson and Steven Feiner. Opportunistic tangible user interfaces for augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 16(1):4–16, 2010.

- [52] Matthew J Hertenstein, Rachel Holmes, Margaret McCullough, and Dacher Keltner. The communication of emotion via touch. *Emotion*, 9(4):566, 2009.
- [53] Ronan Hinchet, Velko Vechev, Herbert Shea, and Otmar Hilliges. Dextres: Wearable haptic feedback for grasping in vr via a thin form-factor electrostatic brake. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, UIST '18, page 901–912, New York, NY, USA, 2018. Association for Computing Machinery.
- [54] Alexandra Ion, Robert Kovacs, Oliver S Schneider, Pedro Lopes, and Patrick Baudisch. A demonstration of metamaterial textures. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–4, 2018.
- [55] Alexandra Ion, Robert Kovacs, Oliver S Schneider, Pedro Lopes, and Patrick Baudisch. Metamaterial textures. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2018.
- [56] Hiroshi Ishii. Tangible bits: Beyond pixels. In *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction*, TEI '08, page xv–xxv, New York, NY, USA, 2008. Association for Computing Machinery.
- [57] Hiroshi Ishii. Tangible bits: beyond pixels. In *Proceedings of the 2nd international conference on Tangible and embedded interaction*, pages xv–xxv, 2008.
- [58] Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labrune. Radical atoms: Beyond tangible bits, toward transformable materials. *Interactions*, 19(1):38–51, jan 2012.
- [59] Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, page 234–241, New York, NY, USA, 1997. Association for Computing Machinery.
- [60] Ali Israr and Freddy Abnoui. Towards pleasant touch: Vibrotactile grids for social touch interactions. In *Extended Abstracts of the 2018 CHI Conference on Human*

- Factors in Computing Systems*, CHI EA '18, page 1–6, New York, NY, USA, 2018. Association for Computing Machinery.
- [61] Alec Jacobson, Daniele Panozzo, et al. libigl: A simple C++ geometry processing library, 2018. <https://libigl.github.io/>.
 - [62] JASP Team. JASP (Version 0.16.3)[Computer software], 2022.
 - [63] Kenneth O Johnson. The roles and functions of cutaneous mechanoreceptors. *Current Opinion in Neurobiology*, 11(4):455–461, 2001.
 - [64] Lynette A Jones, Jacquelyn Kunkel, and Erin Piateski. Vibrotactile pattern recognition on the arm and back. *Perception*, 38(1):52–68, 2009.
 - [65] Hiroyuki Kajimoto, Naoki Kawakami, and Susumu Tachi. Psychophysical evaluation of receptor selectivity in electro-tactile display. In *Proc. of 13th International Symposium on Measurement and Control in Robotics (ISMCR)*, volume 13, pages 83–86, 2003.
 - [66] Smita Kanjanapas, Cara M Nunez, Sophia R Williams, Allison M Okamura, and Ming Luo. Design and analysis of pneumatic 2-dof soft haptic devices for shear display. *IEEE Robotics and Automation Letters*, 4(2):1365–1371, 2019.
 - [67] Mohamed Kari, Tobias Grosse-Puppenthal, Luis Falconeri Coelho, Andreas Rene Fender, David Bethge, Reinhard Schütte, and Christian Holz. TransforMR: Pose-Aware object substitution for composing alternate mixed realities. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 69–79, October 2021.
 - [68] Francisco Kiss, Robin Boldt, Bastian Pfleging, and Stefan Schneegass. Navigation systems for motorcyclists: Exploring wearable tactile feedback for route guidance in the real world. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–7, New York, NY, USA, 2018. Association for Computing Machinery.

- [69] Roberta L. Klatzky, Dianne Pawluk, and Angelika Peer. Haptic perception of material properties and implications for applications. *Proceedings of the IEEE*, 101(9):2081–2092, 2013.
- [70] Katherine J Kuchenbecker, Joseph Romano, and William McMahan. Haptography: Capturing and recreating the rich feel of real surfaces. In *Robotics Research: The 14th International Symposium ISRR*, pages 245–260. Springer, 2011.
- [71] Shinobu Kuroki, Hiroyuki Kajimoto, Hideaki Nii, Naoki Kawakami, and Susumu Tachi. Proposal for tactile sense presentation that combines electrical and mechanical stimulus. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC’07)*, pages 121–126. IEEE, 2007.
- [72] Wallace Santos Lages and Doug A. Bowman. Walking with adaptive augmented reality workspaces: design and usage patterns. In *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI 2019, Marina del Ray, CA, USA, March 17-20, 2019*, pages 356–366, New York, NY, 2019. ACM.
- [73] Susan J Lederman. Tactile roughness of grooved surfaces: The touching process and effects of macro-and microsurface structure. *Perception & Psychophysics*, 16(2):385–395, 1974.
- [74] Susan J Lederman and Roberta L Klatzky. Multisensory texture perception. 2004.
- [75] Edouard Leroy, Ronan Hinchet, and Herbert Shea. Multimode hydraulically amplified electrostatic actuators for wearable haptics. *Advanced Materials*, 32(36):2002564, 2020.
- [76] Vincent Lévesque, Jérôme Pasquero, Vincent Hayward, and Maryse Legault. Display of virtual braille dots by lateral skin deformation: Feasibility study. *ACM Trans. Appl. Percept.*, 2(2):132–149, apr 2005.

- [77] Bruno Lévy, Sylvain Petitjean, Nicolas Ray, and Jérôme Maillot. Least squares conformal maps for automatic texture atlas generation. *ACM transactions on graphics*, 21(3):362–371, July 2002.
- [78] Hongnan Lin, Liang He, Fangli Song, Yifan Li, Tingyu Cheng, Clement Zheng, Wei Wang, and HyunJoo Oh. Flexhaptics: A design method for passive haptic inputs using planar compliant structures. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2022.
- [79] Robert W Lindeman, Yasuyuki Yanagida, Haruo Noma, and Kenichi Hosaka. Wearable vibrotactile systems for virtual contact and information display. *Virtual Reality*, 9:203–213, 2006.
- [80] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. Context-Aware online adaptation of mixed reality interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST ’19, pages 147–160, New York, NY, USA, 2019. Association for Computing Machinery.
- [81] David Lindlbauer and Andrew D Wilson. Remixed reality: manipulating space and time in augmented reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.
- [82] Rong Liu and Hao Zhang. Segmentation of 3D meshes through spectral clustering. In *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings.*, pages 298–305, October 2004.
- [83] Pedro Lopes, Alexandra Ion, and Patrick Baudisch. Impacto: Simulating physical impact by combining tactile stimulation with electrical muscle stimulation. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST ’15, page 11–19, New York, NY, USA, 2015. Association for Computing Machinery.
- [84] Feiyu Lu, Shakiba Davari, and Doug Bowman. Exploration of techniques for rapid activation of glanceable information in head-worn augmented reality. In *Proceedings*

of the 2021 ACM Symposium on Spatial User Interaction, SUI '21, New York, NY, USA, 2021. Association for Computing Machinery.

- [85] Feiyu Lu, Shakiba Davari, Lee Lisle, Yuan Li, and Doug A. Bowman. Glanceable ar: Evaluating information access methods for head-worn augmented reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 930–939, 2020.
- [86] Joseph Luk, Jerome Pasquero, Shannon Little, Karon MacLean, Vincent Levesque, and Vincent Hayward. A role for haptics in mobile interaction: Initial design using a handheld tactile display prototype. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, page 171–180, New York, NY, USA, 2006. Association for Computing Machinery.
- [87] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachsel. Where should we put it? layout and placement strategies of documents in augmented reality for collaborative sensemaking. In *CHI '22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April 2022 - 5 May 2022*, pages 627:1–627:16, New York, NY, 2022. ACM.
- [88] Granit Luzhnica and Eduardo Veas. Investigating interactions for text recognition using a vibrotactile wearable display. *IUI '18*, page 453–465, New York, NY, USA, 2018. Association for Computing Machinery.
- [89] Yuewen Ma, Jianmin Zheng, and Jian Xie. Foldover-Free mesh warping for constrained texture mapping. *IEEE transactions on visualization and computer graphics*, 21(3):375–388, March 2015.
- [90] Jérôme Maillot, Hussein Yahia, and Anne Verroust. Interactive texture mapping. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '93, pages 27–34, New York, NY, USA, September 1993. Association for Computing Machinery.

- [91] Michael R. Marner, Ross T. Smith, James A. Walsh, and Bruce H. Thomas. Spatial user interfaces for large-scale projector-based augmented reality. *IEEE Computer Graphics and Applications*, 34(6):74–82, 2014.
- [92] Fabrice Matulic, Wolfgang Büschel, Michael Ying Yang, Stephan Ihrke, Anmol Ramraika, Carsten Rother, and Raimund Dachsel. Smart ubiquitous projection: Discovering surfaces for the projection of adaptive content. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016, Extended Abstracts*, pages 2592–2600, New York, NY, 2016. ACM.
- [93] Francis McGlone, Johan Wessberg, and Håkan Olausson. Discriminative and affective touch: sensing and feeling. *Neuron*, 82(4):737–755, 2014.
- [94] Ken Nakagaki, Jordan L Tappa, Yi Zheng, Jack Forman, Joanne Leong, Sven Koenig, and Hiroshi Ishii. (dis)appearables: A concept and method for actuated tangible uis to appear and disappear based on stages. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, CHI '22*, New York, NY, USA, 2022. Association for Computing Machinery.
- [95] Ken Nakagaki, Luke Vink, Jared Counts, Daniel Windham, Daniel Leithinger, Sean Follmer, and Hiroshi Ishii. Materiable: Rendering dynamic material properties in response to direct physical touch with shape changing interfaces. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 2764–2772, 2016.
- [96] Benjamin Nuernberger, Eyal Ofek, Hrvoje Benko, and Andrew D. Wilson. Snap-toreality: Aligning augmented reality to the real world. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016*, pages 1233–1244, New York, NY, 2016. ACM.
- [97] A Oulasvirta, N R Dayama, M Shiripour, M John, and A Karrenbauer. Combinatorial optimization of graphical user interface designs. *Proceedings of the IEEE*, 108(3):434–464, March 2020.
- [98] Antti Oulasvirta. User interface design with combinatorial optimization, 2017.

- [99] Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes. *Computational Interaction*. Oxford University Press, 2018.
- [100] Michel Paré, Heather Carnahan, and Allan M Smith. Magnitude estimation of tangential force applied to the fingerpad. *Experimental Brain Research*, 142:342–348, 2002.
- [101] Seonwook Park, Christoph Gebhardt, Roman Rädle, Anna Maria Feit, Hana Vrzakova, Niraj Ramesh Dayama, Hui-Shyong Yeo, Clemens Nylandsted Klokmose, Aaron Quigley, Antti Oulasvirta, and Otmar Hilliges. Adam: Adapting multi-user interfaces for collaborative environments in real-time. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21-26, 2018*, page 184, New York, NY, 2018. ACM.
- [102] Fabrizio Pece, Juan Jose Zarate, Velko Vechev, Nadine Besse, Olexandr Gudozhnik, Herbert Shea, and Otmar Hilliges. Magtics: Flexible and thin form factor magnetic actuators for dynamic and wearable haptic feedback. In *Proceedings of the 30th annual ACM symposium on User interface software and technology*, pages 143–154, 2017.
- [103] Ken Pfeuffer, Yasmeeen Abdrabou, Augusto Esteves, Radiah Rivu, Yomna Abdelrahman, Stefanie Meitner, Amr Saadi, and Florian Alt. ARtention: A design space for gaze-adaptive user interfaces in augmented reality. *Computers & graphics*, 95:1–12, April 2021.
- [104] Domenico Prattichizzo, Claudio Pacchierotti, Stefano Cenci, Kouta Minamizawa, and Giulio Rosati. Using a fingertip tactile device to substitute kinesthetic feedback in haptic interaction. In *Haptics: Generating and Perceiving Tangible Sensations: International Conference, EuroHaptics 2010, Amsterdam, July 8-10, 2010. Proceedings, Part I*, pages 125–130. Springer, 2010.
- [105] Ramesh Raskar and Kok-Lim Low. Interacting with spatially augmented reality. AFRIGRAPH '01, page 101–108, New York, NY, USA, 2001. Association for Computing Machinery.

- [106] Pedro V Sander, John Snyder, Steven J Gortler, and Hugues Hoppe. Texture mapping progressive meshes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 409–416, New York, NY, USA, August 2001. Association for Computing Machinery.
- [107] H Seo and F Cordier. Constrained texture mapping using image warping. *Computer graphics forum: journal of the European Association for Computer Graphics*, 29(1):160–174, March 2010.
- [108] Adalberto L Simeone, Eduardo Velloso, and Hans Gellersen. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 3307–3316, New York, NY, USA, April 2015. Association for Computing Machinery.
- [109] Hyunyoung Song, François Guimbretière, Chang Hu, and Hod Lipson. Modelcraft: Capturing freehand annotations and edits on physical 3d models. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, UIST '06, page 13–22, New York, NY, USA, 2006. Association for Computing Machinery.
- [110] Misha Sra, Sergio Garrido-Jurado, and Pattie Maes. Oasis: Procedurally generated social virtual spaces from 3D scanned real spaces. *IEEE transactions on visualization and computer graphics*, 24(12):3174–3187, December 2018.
- [111] Mandayam A Srinivasan and Robert H LaMotte. Tactual discrimination of softness. *Journal of neurophysiology*, 73(1):88–101, 1995.
- [112] Evan Strasnick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. Haptic links: Bimanual haptics for virtual reality using variable stiffness actuation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–12, New York, NY, USA, 2018. Association for Computing Machinery.
- [113] Keisuke Suzuki, Sohei Wakisaka, and Naotaka Fujii. Substitutional reality system: A novel experimental platform for experiencing alternative reality. *Scientific reports*, 2(1):1–9, June 2012.

- [114] Tomu Tahara, Takashi Seno, Gaku Narita, and Tomoya Ishikawa. Retargetable AR: Context-aware augmented reality in indoor scenes based on 3D scene graph. August 2020.
- [115] Yujie Tao, Shan-Yuan Teng, and Pedro Lopes. Altering perceived softness of real rigid objects by restricting fingerpad deformation. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, page 985–996, New York, NY, USA, 2021. Association for Computing Machinery.
- [116] Wouter M Bergmann Tiest. Tactual perception of material properties. *Vision research*, 50(24):2775–2782, 2010.
- [117] Wouter M Bergmann Tiest and Astrid ML Kappers. Cues for haptic perception of compliance. *IEEE transactions on haptics*, 2(4):189–199, 2009.
- [118] Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. Hapticprint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th annual ACM symposium on user interface software & technology*, pages 583–591, 2015.
- [119] Brygg Ullmer and Hiroshi Ishii. Emerging frameworks for tangible user interfaces. *IBM systems journal*, 39(3.4):915–931, 2000.
- [120] T Aisling Whitaker, Cristina Simões-Franklin, and Fiona N Newell. Vision and touch: Independent or integrated systems for the perception of texture? *Brain research*, 1242:59–72, 2008.
- [121] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–12, New York, NY, USA, 2018. Association for Computing Machinery.
- [122] Anusha Withana, Daniel Groeger, and Jürgen Steimle. Tacttoo: A thin and feel-through tattoo for on-skin tactile output. In *Proceedings of the 31st Annual ACM*

- Symposium on User Interface Software and Technology*, UIST '18, page 365–378, New York, NY, USA, 2018. Association for Computing Machinery.
- [123] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, page 143–146, New York, NY, USA, 2011. Association for Computing Machinery.
 - [124] Jason Wu, Xiaoyi Zhang, Jeff Nichols, and Jeffrey P Bigham. Screen parsing: Towards reverse engineering of ui models from screenshots. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, page 470–483, New York, NY, USA, 2021. Association for Computing Machinery.
 - [125] Xiang, Xu, Hasan, and others. Neutex: Neural texture mapping for volumetric neural rendering. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
 - [126] Chang Xu, Yuxiang Wang, Steven C Hauser, and Gregory J Gerling. In the tactile discrimination of compliance, perceptual cues in addition to contact area are required. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 62, pages 1535–1539. SAGE Publications Sage CA: Los Angeles, CA, 2018.
 - [127] Ya-Ting Yue, Yong-Liang Yang, Gang Ren, and Wenping Wang. SceneCtrl: Mixed reality enhancement via efficient scene editing. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, pages 427–436, New York, NY, USA, October 2017. Association for Computing Machinery.
 - [128] Clement Zheng, Zhen Zhou Yong, Hongnan Lin, HyunJoo Oh, and Ching Chiuan Yen. Shape-haptics: Planar & passive force feedback mechanisms for physical interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2022.